

This Page Is Inserted by IFW Operations
and is not a part of the Official Record

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images may include (but are not limited to):

- BLACK BORDERS
- TEXT CUT OFF AT TOP, BOTTOM OR SIDES
- FADED TEXT
- ILLEGIBLE TEXT
- SKEWED/SLANTED IMAGES
- COLORED PHOTOS
- BLACK OR VERY BLACK AND WHITE DARK PHOTOS
- GRAY SCALE DOCUMENTS

IMAGES ARE BEST AVAILABLE COPY.

**As rescanning documents *will not* correct images,
please do not report the images to the
Image Problem Mailbox.**

This Page Blank (uspto)



The
Patent
Office

PC/GB94/03229



INVESTOR IN PEOPLE

**PRIORITY
DOCUMENT**
SUBMITTED OR TRANSMITTED IN
COMPLIANCE WITH RULE 17.1(a) OR (b)

GB99
3299

The Patent Office
Concept House
Cardiff Road
Newport
South Wales
NP10 8QQ

REC'D 23 NOV 1999

WIPO PCT

I, the undersigned, being an officer duly authorised in accordance with Section 74(1) and (4) of the Deregulation & Contracting Out Act 1994, to sign and issue certificates on behalf of the Comptroller-General, hereby certify that annexed hereto is a true copy of the documents as originally filed in connection with the patent application identified therein.

In accordance with the Patents (Companies Re-registration) Rules 1982, if a company named in this certificate and any accompanying documents has re-registered under the Companies Act 1980 with the same name as that with which it was registered immediately before re-registration save for the substitution as, or inclusion as, the last part of the name of the words "public limited company" or their equivalents in Welsh, references to the name of the company in this certificate and any accompanying documents shall be treated as references to the name with which it is so re-registered.

In accordance with the rules, the words "public limited company" may be replaced by p.l.c., plc, P.L.C. or PLC.

Re-registration under the Companies Act does not constitute a new legal entity but merely subjects the company to certain additional company law rules.


Signed

Dated

-3 NOV 1999

This Page Blank (uspto)

Patents Form 1/77
Patents Act 1977
(Rule 16)



**The
Patent
Office**

07JUL99 E460201-2 D02917
P01/7700 0.00 - 9915838.8

Request for grant of a patent

The Patent Office
Cardiff Road
Newport
Gwent NP9 1RH

1.	Your reference 5279302/EMR	
2.	Patent Application Number	9915838.8 - 6 JUL 1999
3.	Full name, address and postcode of the or of each applicant (<i>underline all surnames</i>) University College London Gower Street London WC1E 6BT Patents ADP number (<i>if known</i>) 798652002 If the applicant is a corporate body, give the country/state of its incorporation Country: ENGLAND State:	
4.	Title of the invention ENERGY PROPAGATION MODELLING APPARATUS	
5.	Name of agent "Address for Service" in the United Kingdom to which all correspondence should be sent Patents ADP number	Beresford & Co 2/5 Warwick Court High Holborn London WC1R 5DJ 1826001
6.	Priority details	
	Country	Priority application number
	GB	9821139.4
		Date of filing
		29 September 1998

Patents Form 1/77

7. If this application is divided or otherwise derived from an earlier UK application give details

Number of earlier of application

Date of filing

8. Is a statement of inventorship and or right to grant of a patent required in support of this request?

YES

9. Enter the number of sheets for any of the following items you are filing with this form.

0 Continuation sheets of this form

73 Description

11 Claim(s)

1 Abstract

32* 32 Drawing(s)

10. If you are also filing any of the following, state how many against each item.

0 Priority documents

0 Translations of priority documents

1 Statement of inventorship and right to grant of a patent (*Patents form 7/77*)

0 Request for preliminary examination and search (*Patents Form 9/77*)

0 Request for Substantive Examination (*Patents Form 10/77*)

1 Any other documents (*please specify*)

Form 23/77

11. I/We request the grant of a patent on the basis of this application

Signature


BERESFORD & Co

Date

6 July 1999

12. Name and daytime telephone number of person to contact in the United Kingdom

ROUND; Edward Mark

Tel:0171-831-2290

ENERGY PROPAGATION MODELLING APPARATUS

The present invention is concerned with apparatus for representing the modelling of the propagation of energy.

5 In particular, the present invention is concerned with the modelling of propagation which can be represented by ray paths, and the manner in which energy passes along ray paths while interacting with objects in a modelled scene. The present invention has applications in the
10 modelling of light in a scene, to provide a facility for production of near photo-realistic images of a globally illuminated scene or to provide for measurement of light or other energy distribution within a scene.

15 In many computer games, dynamic imagery is of real importance, and so real-time production of dynamic images is implemented to the detriment of photo-realism. On the other hand, many applications such as architectural design do not require substantial dynamic changes to an
20 image in real-time, whereas photo-realism is of real importance.

However, even in a photo-realistic model, it is preferable that a user can change the viewpoint of a
25 computer modelled scene, with limited re-calculation. An ideal technique, for carrying this out without significant re-calculation, is known as a "view-

independent" technique. View-independence in this sense means that data pertaining to the scene does not require new lighting calculations on variation of the point of view of the scene.

5

In order to model a scene in a photo-realistic manner, the true reflective and transmissive properties of light incident on objects in the scene must be approximated. In particular, the reflection and transmission of light must be modelled so as to approximate to the behaviour of light in real life. In that way, the observer of the modelled scene will reach the correct conclusions concerning the nature and appearance of objects within the scene.

15

Light within a scene has in the past been modelled by several techniques, most of which can be categorised as either "ray tracing" techniques or "radiosity" techniques.

20

Several problems are associated with ray tracing. Firstly, ray tracing is view-dependent. That means that if a different viewpoint is selected, the transmission of light towards the viewpoint from the scene must be entirely re-calculated. Moreover, it is not suitable for modelling diffuse reflection, where a particular incident ray results in a plurality of reflected ray directions,

and a particular reflected ray can be the result of a plurality of incident rays. It is impossible to represent diffuse reflection realistically using ray tracing.

5

Several modifications of ray tracing have been proposed, but with many drawbacks. In several cases, diffuse surfaces can be considered, alongside specular reflecting and transmitting surfaces. In one such technique

10 diffusely reflecting surfaces are treated using discrete elements thereof, and specularly reflecting surfaces are taken into account. However, subsequent rendering results in specularly reflecting surfaces not being displayed as such.

15

Radiosity techniques were introduced in "Modelling the interaction of light between diffuse surfaces" by C M Goral et al, published in Computer Graphics (SIGGRAPH 84), 213-222. Radiosity is defined as the rate at which

20 light energy leaves a unit area of a surface. Radiosity techniques are suitable for scenes where all materials are ideal diffuse reflectors. However, radiosity is not suitable for consideration of specular reflection.

25

The lack of specular reflective effects means that the lighting of a model considered through radiosity techniques is independent of the position of the

viewpoint of the model. That is, recalculation of the lighting need not take place in the event that the position of the viewpoint changes. This is described as "view independence" in an alternative sense of the term.

5

As a result, a radiosity model is suitable to be viewed from a dynamically changing viewpoint in real time. This is known as "real-time walk-through". This is highly suitable for CAD applications such as for architectural models, but the loss of specular reflection results in a loss in photo-realism.

10

Finally, a rendering technique proposed on the one-hand by Levoy and Hanrahan in a paper published in Computer Graphics (SIGGRAPH), Annual Conference Series (1996), 31-42, and on the other-hand by Gortler et al in Computer Graphics (SIGGRAPH), Annual Conference Series (1996), 43-52, does not fall within the scope of either ray tracing or radiosity methods, but may make use of aspects of both.

20

In both disclosures, a system implementing a "light field" or "lumigraph" is described. The light field is constructed of an ordered plurality of rays constructed between two parallel planes. Other light fields can be constructed using different arrangements of construction planes. The rays are arranged between discrete points

25

on the planes, and can be described in terms of the points on the two planes between which they extend. A scene consisting of one or more objects can be rendered into the light field, by considering views of the scene at each point on the two planes. Those views can either be constructed by one of the techniques identified above, or by importing real views of a scene produced on digital cameras. The radiance propagated along each ray from a particular point constructing the light field can be identified from the view from that point constructed or imported by the system.

Once the radiances associated with rays from each point have been identified, views of the scene from a point within the light field can be constructed by considering the radiance of rays passing near that point and interpolating an image from the radiance values obtained.

Real images of real objects can be rendered into the light field, so that images from positions other than the positions from which images were obtained by a camera can also be seen. These real images can be generated by means of a digital camera, taking digital photographs from a plurality of different positions about the scene to be rendered into the light field.

However, the light field technique described above has

a significant disadvantage, in that views cannot be constructed from every position in the light field. If a scene consists of a plurality of objects, or at least one object which has at least one concave surface, then

5 a volume can be defined, between a surface bounding a minimum volume about the scene which includes no concave surfaces (the "convex hull" of the scene) and the actual surface(s) defining the object(s) of the scene. Several rays within this volume are bounded at both ends by

10 surfaces of objects in the scene. Therefore, radiances along those rays cannot be determined from consideration of real images at points on the two constructing planes. As a result, discontinuities exist which prevent construction of images from a sub-set of all viewpoints

15 within the volume.

It is an object of the present invention to provide an alternative technique for modelling energy propagation, such as the propagation of light through a scene. The

20 invention provides simulation apparatus which defines discrete paths within an environment within which energy propagation is to be modelled, and which locates points along those paths which denote interactions of objects within the environment with the paths.

25

One aspect of the present invention provides a technique which allows specular and diffuse reflection to be

represented realistically. A scene, and light propagation within the scene, can be represented by the present invention such that an image of the scene from a particular viewpoint can be generated in substantially constant time. The constant time is not dependent on the complexity of the scene to be represented and, with the provision of sufficient computer power, is capable of being made significantly small that real time changes of viewpoint can be represented.

10

Further features and advantages of the invention will become apparent from the following description of apparatus and a method in accordance with a preferred and specific embodiment of the invention, with reference to the accompanying drawings in which:

15

Figure 1 is a schematic diagram of image processing apparatus in accordance with a specific embodiment of the invention;

20

Figure 2A is a perspective view of a scene to be modelled by the image processing apparatus of Figure 1, including illustration of ray directions encountered at a first viewpoint;

25

Figure 2B is a perspective view of the scene illustrated in Figure 2A, including illustration of ray directions

encountered at a second viewpoint;

Figure 3A is a perspective view of a parallel sub-field
in accordance with a specific embodiment of the present
5 invention;

Figure 3B is a perspective view of a further parallel
sub-field in accordance with a specific embodiment of the
present invention;

10

Figure 3C is a perspective view of a still further
parallel sub-field in accordance with a specific
embodiment of the present invention;

15 Figure 4 is a perspective view of the scene illustrated
in Figures 2A and 2B, including illustration of a
selected plane within the scene to be modelled and a
selected ray within that plane in accordance with the
specific embodiment of the present invention;

20

Figure 5 is an elevation of the construction plane
illustrated in Figure 4 in a direction normal to the
construction plane, in accordance with the specific
embodiment of the present invention;

25

Figure 6 is a schematic diagram illustrating
intersections of objects with the light ray illustrated

in Figure 4, throughout its length;

Figure 7 is a schematic diagram showing an extract of the length of the light ray, between the side walls of the scene illustrated in Figure 4;

Figure 8 is a schematic diagram showing the structure of interval data in accordance with the specific embodiment of the present invention;

10

Figure 9 is a schematic diagram of a data structure for the interval data illustrated in Figure 8;

Figure 10 is an elevation similar to that illustrated in Figure 5, including a representation of a virtual eye for viewing the scene;

Figure 11 is a schematic diagram showing the structure of the image processing apparatus illustrated in Figure 1;

20

Figure 12 is a schematic diagram showing the structure of the L F computation unit illustrated in Figure 11;

Figure 13 is a schematic diagram showing the internal structure of the viewer illustrated in Figure 11;

25

Figure 14 is a flow diagram illustrating the procedure performed by the image processing apparatus in use;

Figures 15 to 20 are flow diagrams illustrating sub-procedures called by the procedure illustrated in Figure 14;

Figure 21 is a flow diagram illustrating a procedure according to which the viewer operates in use;

10

Figure 22 is a flow diagram illustrating a procedure by which the image processing apparatus is operative to delete an object from a scene;

15 Figure 23 is a flow diagram illustrating a procedure by which the image processing apparatus is operative to add an object to a scene;

Figure 24 is a schematic diagram of an object in a light field demonstrating the potential for computational complexity in considering diffuse reflection;

Figure 25 is a schematic diagram showing a portion of the surface of the object illustrated in Figure 24, illustrating a gathering step of a method of reducing computational complexity;

Figure 26 is a diagram of the portion of the object surface illustrated in Figure 25, demonstrating a "shooting" step of the method of reducing computational complexity;

5

Figure 27 is a view of the object illustrated in Figure 24, demonstrating a further method of reducing the computational complexity of considering diffuse reflection;

10

Figure 28 is a schematic diagram of image processing apparatus for defining an environment in virtual reality, in accordance with the specific embodiment of the invention;

15

Figure 29 is a side view of a head mounted display for use in the apparatus illustrated in Figure 28; and

20

Figure 30 is a view of the display apparatus of the head mounted display illustrated in Figure 29.

25

Figure 1 is a block diagram showing the general arrangement of an image processing apparatus according to an embodiment of the invention. In the apparatus, there is provided a computer 2, which comprises a central processing unit (CPU) 4 connected to a memory 6 operable to store a programme defining the sequence of operations

of the CPU 4 and to store object and image data used in calculations by the CPU 4.

5 An input device 8 is coupled to an input port (not shown) of the CPU 4. The input device 8 may comprise, for example, a keyboard and/or a position sensitive input device such as a mouse, tracker ball, or a digitizer tablet and stylus etc.

10 A frame buffer 10 is also coupled to the CPU 4, the frame buffer 10 comprising a memory unit (not shown) arranged to store image data relating to at least one image, for example by providing one (or several) memory location(s) per pixel of the image. The value(s) stored in the frame
15 buffer 10 for each pixel defines the colour or intensity of that pixel in the image.

In the present embodiment of the invention, an image is represented as a two-dimensional array of pixels, which
20 can conveniently be described in terms of Cartesian co-ordinates. The position of the given pixel can be described by a pair of x, y co-ordinates. The frame buffer 10 has sufficient memory capacity to store at least one image. If the image has a resolution of 1000
25 by 1000 pixels, the frame buffer 10 should include 10^6 pixel locations, each location being addressable directly or indirectly in terms of pixel co-ordinates x,y.

A video display unit (VDU) 12 is coupled to the frame buffer 10. The VDU 12 is operative to display the image stored in the frame buffer 10 in a conventional manner. For instance, if the VDU 12 displays images in a raster scanning manner, the x co-ordinate of a pixel maps to the distance along a line of the scanned display, and the y co-ordinate of the pixel maps to the number of the line.

Also coupled to the frame buffer 10 is a video tape recorder (VTR) 14 adapted to receive a video tape 15. Another image recording device, such as a paper printer a 35mm film recorder or a recordable compact disc could be provided in addition or in the alternative.

A mass storage device 16, such as a hard-disk drive, is coupled to the memory 6. The mass storage device 16 has a high data storage capacity, suitable for storing data to which instant access is not required. Moreover, a disk drive 18, operable to accept removable data storage media such as a floppy disk 20 is coupled to the memory 6. The disk drive 18 is operable to transfer data stored on the floppy disk 20 to the memory 6.

Figure 2A illustrates a scene 22 in respect of which it would be desirable to model light propagation using the apparatus of the embodiment of the invention described herein. The scene 22 consists of a room bounded by left

and right side walls 24, 26, back and front walls 28, 30, a floor 32 and a ceiling 34. The front wall 30 and the ceiling 34 are shown as transparent in Figure 2A for reasons of clarity. The walls, floor and ceiling
5 together define a room which, in this embodiment, is substantially cubic in shape.

A substantially square window 36 is set in the left-hand side wall 24, with edges parallel to the edges of the
10 wall 24. A rectangular mirror 38 is mounted on the back wall 28 with the longer sides of the rectangular mirror 38 being substantially vertical.

A standard lamp 40 of conventional size stands on the
15 floor 32, substantially in the corner defined by the left-hand side wall 24 and the back wall 28. A chair 42 is situated substantially in the centre of the floor 32. The chair 42 is of simple construction, having a seat, a back and four legs. The chair 42 is oriented in the
20 room such that its back is substantially parallel with a line constructed between the back left corner of the room (in which the standard lamp 40 stands) and the front right corner of the room. The chair 42 faces the front left corner of the room.

25

In Figure 2A, an eye 44 is placed at a viewpoint at the front of the scene 22 and slightly to the left of centre.

As is well-known from conventional optics, light enters the eye 44 and is refracted by various elements of the eye 44, to allow a focused image to be formed on the retina.

5

In the scene 22, rays of light travel in an infinite number of directions. However, only those rays which enter the pupil of the eye 44 are experienced by the observer. For example, a ray 46 is shown travelling from the window 36 (transmitting light from a theoretically infinitely distant light source), directly to the eye 44. Another ray 48 is shown travelling directly from the standard lamp 40 to the eye 44.

15 However, a substantial amount of light will not reach the eye 44 directly from light sources, but will be reflected off objects within the scene 22. For example, light radiated from the standard lamp 40 will hit the left-hand wall 24, the floor 32 and the right-hand wall 26, and be
20 reflected towards the eye 44. Rays 50a, 50b, 50c respectively illustrate these three sets of circumstances. It will be understood that light radiated from the lamp will also be reflected towards the eye from the back wall 28, the front wall 32 and the ceiling 34,
25 but rays corresponding to those paths are omitted from Figure 2A for reasons of clarity.

Moreover, rays 52a, 52b in Figure 2A represent light travelling from the window 36 to the eye 44 via a reflection on the back of the chair 42. Rays 54a, 54b represent light travelling from the lamp 40 to the eye 44 via a reflection of the seat of the chair 42.

The walls, floor and ceiling 24-34 and the chair 42 have diffusely reflecting surfaces. Consequently, a ray of light incident on one of those objects will reflect in all directions forming an acute angle with the surface normal at the point of incidence. Unit vectors representing those directions define a hemispherical volume of unit radius, with centre at the point of incidence, and bounded by a plane tangential to the surface of the object at the point of incidence.

As a result, an observer of light reflected at one of the objects with diffusely reflective surfaces will see a representation of the object in question. The observer will not see a sharp reflected image of the ultimate source of the reflected light. Also, some of the light incident on each of these objects will be absorbed thereby. The rate of absorption is normally frequency dependent, and the absorption is associated with interpretation by the observer of colour of the object in question.

In contrast, the ray 56 illustrated in Figure 2A, which is firstly diffusely reflected from the back of the chair 42 towards the mirror 38, is then specularly reflected from the mirror towards the eye 44. Specular reflection is that which behaves in accordance with the ideal laws of reflection, as described in the introduction above. By observing other rays which similarly diffusely reflect from the various objects in the scene to the mirror and then to the eye, a reflected image of the scene is formed in the eye.

Referring now to Figure 2B, the same scene 22 is illustrated, with the eye of the observer in a new position. The eye is now referenced with the numeral 44'. The objects within the scene 22 described above are given the same reference numerals as before, in view of the fact that they are unchanged. However, the scene 22 is now being viewed from a different position, and so different rays of light will enter the eye 44'. For example, a ray 58 travels directly from the window 36 to the eye 44', and a ray 60 travels directly from the lamp 40 to the eye 44'. Furthermore, rays 62a, 62b travel from the lamp 40 via diffuse reflection of the chair 42 to enter the eye 44'. A ray 64 will travel from the lamp 40, diffusely reflect from the back wall 28, and enter the eye 44'. A ray 66 from the lamp 40 will specularly reflect from the mirror 38 and enter the eye 44'. That

ray 66 is representative of rays from the lamp which reflect from the mirror 38, which will contribute to the construction of a reflected image of the lamp within the eye.

5

From the foregoing description of the two views of the same scene, it will be seen that in order to observe the same scene from two different views, one must obtain information concerning the radiance of light associated with two different sets of rays.

10

In the real world, light exists independently of a viewer, and behaves in accordance with physical laws such as concerning reflection, refraction and diffraction.

15

Where a view of a scene is observed in the real world, those rays which enter the aperture of a viewer are taken into consideration, and focused to produce an image. The preferred embodiment of the present invention aims to emulate that situation, by modelling a complete pattern of light behaviour within a scene, without reference to a particular viewpoint.

20

This objective is not achievable in practice, because any scene contains an infinite number of potential light propagation directions. Therefore, the present invention as illustrated by the preferred embodiment models a light field consisting of a finite number of potential carriers

25

of light radiance, referred to hereinafter as "rays".

Thereafter, objects can be rendered into the light field by considering the intersections of the rays in the light field with the objects. Once all objects have been rendered into the light field, light can be propagated along the rays from objects designated as light emitters, under interruptions such as reflections, refractions and defractions can be computed at the intersections of the rays with the objects.

The preferred embodiment is significantly different from the techniques described by Levoy and Hanrahan and Gortler et al, both of which are directed to the construction of a 3D representation of a scene from 2D images. In contrast, the preferred embodiment described herein is concerned primarily with the behaviour of light within the light field, and considers the characteristics of the objects rendered into the light field only when calculating interactions at intersections.

Series of 2D images need not be stored with the present invention; a final image at a particular viewpoint in the light field can be constructed in a straightforward manner.

Various different light fields could be defined, but a

preferred light field consists of a plurality of parallel sub-fields. Each parallel sub-field consists of a light region, of generally cuboid shape, between two sides of which extend a plurality of parallel rays. The plurality
5 of parallel rays is arranged in a grid to enable straightforward referencing thereof. The centre of each parallel sub-field is defined as its origin, and the parallel sub-fields are overlaid so as to have a common origin. The parallel sub-fields are oriented in
10 different directions in relation to each other, so that a light field is constructed which includes rays offset across a region, and oriented in a variety of directions in three-dimensional space.

15 With reference now to Figure 3A, a parallel sub-field 70 comprises a cubic volume in x, y, z Cartesian space defined by:

$$\begin{aligned} -1 &\leq x \leq 1 \\ 20 \quad -1 &\leq y \leq 1 \\ -1 &\leq z \leq 1 \end{aligned}$$

A rectangular grid is imposed on the xy plane comprising n sub-divisions of equal width in each of the x and y
25 directions. The width of each sub-division in each of the x and y directions is $2/n$.

Discrete co-ordinates can be assigned to the centre of these sub-divisions, imposing a discrete co-ordinate grid on the x, y plane within the parallel sub-field 70. The grid can be referenced by co-ordinates (i, j). The real
 5 co-ordinate on the x, y plane corresponding to (i, j) is therefore:

$$(x_i, y_j) = \left(\frac{2i+1}{n} - 1, \frac{2j+1}{n} - 1 \right); i=0, 1, \dots, n-1; j=0, 1, \dots, n-1$$

Each grid reference has associated therewith a ray 72
 10 parallel with the z axis and spanning the cubic volume. Therefore there are n x n rays, defined in terms of co-ordinates (i, j), each of which coordinates range between 0 and n-1.

15 In that way, a parallel sub-field 70 is defined which contains a plurality of rays 72 in the vertical direction, which can be used to model light travelling in either direction along each ray in that set of rays.

20 With reference to Figure 3B, the cubic volume 70 previously described is illustrated rotated through an angle ϕ from the vertical about the y axis. In that way, each ray of the n x n array of rays 72 is oriented at an angle ϕ from the vertical and when resolved into the x,

y plane is parallel with the x axis.

Figure 3C illustrates the cubic volume containing the array of rays further rotated through an angle θ about the z axis. In that way, each of the rays of the $n \times n$ array is oriented at an angle ϕ from the vertical and, when resolved into the x, y axis, is at an angle θ relative the x axis.

10 All possible directions of rays can be described using the following ranges:

$$\phi = 0 ; \theta : \text{don't care}$$

$$0 < \phi < \pi/2 ; 0 \leq \theta < 2\pi$$

15 $\phi = \pi/2 ; 0 \leq \theta < \pi$

It should be noted that at $\phi = 0$, the direction of the rays 72 is independent of the value of θ , and that at $\phi = \pi/2$, θ need only pass through a half revolution in order to cover all possible directions.

In order to implement the arrangement on a computer, θ and ϕ must be discretised. Two alternative methods of discretising ϕ and θ will now be described.

25

In a first method, R is a predetermined number of subdivisions of 2π , representing the level of resolution

which the apparatus is to use. Then, defining the discretised variables u, v corresponding to continuous variables θ, ϕ :

$$\theta = u \frac{2\pi}{R}, \quad u = 0, 1, \dots, R-1$$

$$\phi = v \frac{2\pi}{R}, \quad v = 0, 1, \dots, \frac{R}{4}$$

5

Therefore, by substituting these equations into the range definitions of ϕ and θ and making u, v the subject thereof, the ranges of values which u and v can take are as follows:

10

$v = 0; \quad u : \text{don't care}$

$v = 1, 2, \dots, (R/4)-1; \quad u = 0, 1, \dots, R-1$

$v = R/4; \quad u = 0, 1, \dots, (R/2)-1$

15 Since the above formulae refer to the quantity $R/4$, R must be an integer multiple of 4 for u and v to be integers.

20 However, by using the same resolution to divide the range of θ regardless of the value of ϕ , the rays generated by the above scheme for values of ϕ near zero are very close together. In contrast, when ϕ is near $\pi/2$, the generated

rays are relatively spaced apart. It would be advantageous to be able to generate a set of rays within the hemisphere defined by the variables $\phi[0, \pi/2]$ and $\theta[0, 2\pi]$, the directions of the rays being equispaced.

5 A second method of discretising θ and ϕ in accordance with a specific embodiment of the invention produces improved spacing of ray directions. In that second method, the discrete values which θ and ϕ can assume are defined as follows:

10

$$\phi = 0, \theta = \text{don't care}$$

and

$$\phi = (\pi/2R)u ; u = 1, 2, \dots R$$

$$\theta = (2\pi/3u)v ; v = 0, 1, \dots, 3u - 1$$

15 where

$$R = 2^{m-1}$$

This sampling scheme is governed by a parameter m . The resolution of ϕ and θ increases with the value of m . ϕ is discretised equally over the range $[0, \pi/2]$. θ is sub-divided over the range $[0, 2\pi]$ to a degree which is a variable depending on the value of u , and consequently on the value of ϕ . If ϕ is small, the number of subdivisions of θ is also small. If ϕ is large, the number of sub-divisions of θ is also large.

25

For example, if $m = 5$, then $R = 31$. Accordingly, there

will be 32 sub-divisions of ϕ over its range. In the lower end of the range of ϕ , for instance if $u = 2$, v will vary between 0 and 5. Therefore, there will be six sub-divisions of θ .

5

Conversely, at the upper end of the range of u , such as when $u = 30$, v will take values in the range 0 to 89. Accordingly, if $u = 30$, there will be 90 sub-divisions of the range of θ .

10

This sampling scheme produces sub-divisions of θ and ϕ which provides an approximation to a uniform distribution of rays.

15 Defining rays over the ranges set out above, a complete light field can be represented by the co-ordinate system (i, j, u, v) . This can be stored in the computer memory 6 as a 4D array. For efficiency of memory, this 4D array can be further represented in flattened form as a 1D
20 array.

A scene can be defined within the light field constructed of the arrays of rays. It is convenient to confine the scene for example to a unit cube space defined by the
25 ranges:

$$- 0.5 \leq x \leq 0.5$$

$$- 0.5 \leq y \leq 0.5$$

$$- 0.5 \leq z \leq 0.5$$

This ensures that any orientation of the cubic volume of the parallel sub-field 70 will fully enclose the scene.

5 That means that no part of the scene is omitted from coverage by rays in any particular direction. In fact, since the cubic volume 70 has smallest (edge) dimension being 2, any scene with longest dimension no greater than 2 can be accommodated within the cubic volume 70.

10

Assuming that the real ray 48 in Figure 2A matches with a ray in the light field. It can be identified by a start point (x_1, y_1, z_1) and a direction (dx, dy, dz) . (θ, ϕ) or (u, v) can then be found using trigonometry on
15 direction (dx, dy, dz) . Matrix rotation of the ray 48 will rotate the ray so that it becomes vertical (parallel with the z axis). Then, the (i, j) co-ordinates can be found from the point of intersection of this transformed ray with the x, y plane and the previously stated
20 definition of i, j in terms of x, y .

In the event that the real ray 48 does not match with a ray 72 of the light field, it is necessary to find a ray within the light field which best approximates the real
25 ray. It will be appreciated that for any accepted measurement of error and for any accepted error bound, values for m and n , the discretisation variables, can be

selected.

Firstly, once the real ray directions has been converted into direction coordinates (θ, ϕ) or (u, v) , the direction coordinates are "rounded" to the nearest
 5 provided value. In the case of (u, v) , this will be the nearest integer. Secondly, once the actual intersection of the real ray, rotated back to the vertical, with the x, y plane has been found, the co-ordinates of the
 10 intersection can be "rounded" to the nearest provided value of (x, y) or (i, j) . In the case of (i, j) , this will be the nearest integer. It is preferable to fit the direction first, since errors introduced by approximating direction will have a greater impact on the eventual
 15 propagation of light in the scene.

In reverse, a given set of co-ordinates (i, j, u, v) defines a ray 72 which may be within the scene. u and v can be used to calculate (θ, ϕ) for that ray 72.
 20 Therefore, $(x_i, y_j, -1) R_y(\phi) R_z(\theta)$ and $(x_j, y_i, +1) R_y(\phi) R_z(\theta)$ give the end points of the ray in real co-ordinates, where R_y and R_z are standard three dimensional rotation matrices about the y and z axes respectively. This is useful for identifying position of a point lying on the
 25 ray within the scene.

The scene 22 of Figures 2A and 2B is illustrated again

in Figure 4. A particular plane 74 within the scene is identified by dash lines; the position of the plane 74 is defined by a horizontal line along the left-hand side wall 24 substantially one third of the height of the wall from the top of the wall, and a horizontal line along the right-hand side wall 26 substantially at the foot of the right-hand side wall. The plane 74 intersects the standard lamp 40 and the back of the chair 42. The selected plane 74 contains a large number of rays, depending on the resolution of the discretisation of ϕ and θ . However, Figure 5 shows a selection of these rays 72 in three different directions. All other rays have been omitted for reasons of clarity. The portion of the scene intersected by the plane 74 is shown, with the window 36, the mirror 38, the standard lamp 40 and the back of the chair 42 illustrated.

One particular ray 86 is shown, which passes along the plane through the standard lamp 40 and the back of the chair 42. This is illustrated in more detail in Figure 6. The ray, when transformed back to the vertical, extends from $z = -1$ to $z = +1$, with the various objects intersected, namely the left-hand side wall 24, the standard lamp 40, the back of the chair 42 and the right-hand side wall 26 placed thereon at the appropriate place. All points along the ray can be assigned a position by referring back to the z axis.

Figure 7 illustrates the segment of the ray between the side walls 24, 26 which is the part of the ray which is of most interest, being as it is part of the scene 22.

5 In order to render an object of the scene into the light field, firstly a ray passing through an object must be computed. Then, when such a ray 86 is computed, all of the intersections of that ray with the object must be identified by position. That position is expressed as
10 a value t in the parametric equation:

$$p(t) = p + t(q-p) \quad 0 \leq t \leq 1$$

where p and q correspond with intersections of the ray
15 86 with the limits of the scene, i.e. side walls 24, 26. Therefore, as intersections are found, the ray 86 has an associated set of intersections $[0, t_1, t_2, \dots, t_k, 1]$ where the t_i are parametric values corresponding to the positions of the intersections with objects. The real
20 co-ordinate position of each intersection can be found from $p(t_i)$.

In order to keep account of intersections of the ray, the intersections are arranged in an interval list. Interval
25 lists for all rays of the entire light field are arranged in a four-dimensional array indexed by parameters $(i, j,$

u, v). Each interval list is initially set to null, before intersections are found and arranged therein. An initial interval list represents the interval $[0, 1]$, i.e. the full length of the ray within the bounds of the scene.

Each entry in the interval list includes not only the t value, but other data concerning the nature of the intersection as well. All of the data is collected in a structure called, hereinafter, a T-Intersection, having fields for the t-value, the object intersected, the direction of potential radiance from the intersection, a radiance r, and an unshot radiance u. Radiance r and unshot radiance u are typically vectors, expressing energy levels for red, green and blue light (RGB) so that coloured light can be represented. These vectors are initially set to zero vectors. It will be appreciated that other means of expressing coloured light, other than RGB, are also possible.

20

The direction is set to be either "right" or "left". The "right" direction corresponds with the direction along the ray from $z < 0$ to $z > 0$ when the ray is transformed back to the vertical; "left" is in the opposite direction.

25

Radiance r and unshot radiance u may be vector quantities if the system is arranged to consider colour,

polarisation or other characteristics of light within the system which need to be described. The structure of the T-Intersections is illustrated in Figure 8. It may also be advantageous to include a field for the "normal vector" at the point of intersection with the object, though this is not essential since the normal can always be computed by reference to the point of intersection and the data defining the object.

10 The T-Intersections should be placed in a data structure which allows for ease of implementation and relatively constant look-up time.

In this embodiment, a standard recursive binary tree is used, known hereinafter as an interval tree. The interval tree corresponding to the ray 86 illustrated in Figure 7 is illustrated in Figure 9.

In the general case, each node of the interval tree corresponds to a T-Intersection and further comprises two child nodes, called left_Tree and right_Tree. left_Tree and right_Tree are also interval trees. When an object is rendered into the light field, and a first intersection is found with the ray under consideration, the intersection is loaded into the previously empty interval tree for that ray. The interval tree now comprises one node, and the two child nodes are null.

As a second T-intersection is loaded into the interval tree, the t-value contained in that second T-Intersection is compared with the t-value contained in the T-Intersection of the first node. If the t-value of the second T-Intersection is lower than the t-value of the first T-Intersection, then the second T-Intersection is placed in the left-Tree child node of the first node. Otherwise, the second T-Intersection is placed in the right Node child node of the first node. In that way, an interval tree is constructed which is sorted in respect of the t-values, which binary tree can be used with ease to search for T-Intersections by t-value.

A third intersection can be added in the same way - if a child node is full then comparison is made with the t-value of the contents of that child node. Progress is made down the branches of the tree until a null child node at an appropriate position is found.

In the example illustrated in Figure 7, the ray 86 is parameterised such that $t = 0$ at the left-hand wall 24 and $t = 1$ at the right-hand wall 26. The T-Intersections along the ray 86 are identified by subscript, with T_1 to T_4 representing intersections with the chair 42 and T_5 and T_6 representing intersections with the lamp 40. T_7 and T_8 represent the intersections with the side walls

24, 26. Accordingly, the T-Intersections have the following attributes:

5 $T_1 = (0.50, 42, \text{left}, 0, 0)$
 $T_2 = (0.52, 42, \text{right}, 0, 0)$
 $T_3 = (0.70, 42, \text{left}, 0, 0)$
 $T_4 = (0.72, 42, \text{right}, 0, 0)$
 $T_5 = (0.07, 40, \text{left}, 0, 0)$
 $T_6 = (0.15, 40, \text{right}, 0, 0)$
10 $T_7 = (0.00, 24, \text{right}, 0, 0)$
 $T_8 = (1.00, 26, \text{left}, 0, 0)$

15 If those T-Intersections are found in that order, and loaded onto a binary tree, the binary tree will take the form illustrated in Figure 9.

20 Once T-Intersections are loaded into the binary tree, look-up time is dependent on the logarithm of the number of T-Intersections. This is an advantageous arrangement because look-up time in the interval tree will not increase significantly as an increasing number of intersections are loaded thereon. Moreover, as will become apparent from later description, at times it is necessary to find a T-Intersection adjacent a given T-Intersection in a given direction. It is relatively straightforward to find that T-Intersection using the

25

binary tree in its conventional way.

Once all rays have been found which intersect objects,
and all intersections have been loaded onto relevant
5 binary trees, the rendering of the scene into the light
field is considered complete.

Thereafter, radiance must be added to the light field,
by taking into consideration any objects within the field
10 which emit light. In the present example, the window 36
and the standard lamp 40 are considered to be light
emitters.

Taking into consideration the standard lamp 40, all of
15 those rays which intersect the standard lamp are
computed. When a ray 86 is computed, a T-Intersection
 T_5 is identified on the ray as intersecting with the
standard lamp 40. Radiance in accordance with the light
emission characteristic of the standard lamp 40 is added
20 to the data structure, and unshot radiance equal to the
change in radiance is also added thereto. T-Intersection
 T_6 is then found and treated in the same way, and
thereafter, all rays intersecting the standard lamp 40
are treated in the same way.

25

Moreover, although not illustrated in Figure 5, rays
intersecting the window 36 are treated in the same way

with respect to light emitted thereby, or more properly transmitted therethrough. In the interests of simplicity, it is more straightforward to consider the window as an object with light emitting properties than
5 to render the sun into the scene.

Obviously, if it is desired to represent further objects on the other side of the window, those objects must be rendered into the light field as well. This may require
10 some scaling of the scene in order to ensure that the whole scene is enclosed.

Once all light emission has been added to the light field, each object is considered to establish whether it
15 is intersected by a ray which carries unshot radiance. In the example shown in Figures 5 and 7, the back of the chair 42 is in receipt of unshot radiance along segment (T_6 , T_1). At that point, the unshot radiance is considered to be emitted light from that point on the
20 chair, and it is transmitted through diffuse reflection along all rays emanating from on or near T-Intersection T_1 on ray 86.

In the same way as a real ray is unlikely to match
25 exactly a ray in the light field, it is unlikely that exact coincidence with other rays will occur at intersection T_1 , and so approximations will be necessary.

In fact, in each desired direction, the closest ray is selected, and the diffusely reflected radiance and unshot radiance is added to each such ray. Those diffuse reflections are identified by arrows 88, 90, 92, 94 in Figure 5. Moreover, some radiance will be reflected back onto ray 86 towards the standard lamp 40. Rays 88, 90 are shown despite the fact that they do not coincide with illustrated sets of rays; they will each coincide with another ray not illustrated but within the set of rays in the light field which are contained in the illustrated plane 74. Other diffusely reflected rays will also be identified in other directions not contained in the plane 74.

Once all unshot radiance incident on the chair 42 is dealt with in this way, the next object to be in receipt of unshot radiance is considered. For example, the mirror 38 is now in receipt of unshot radiance as a result of reflected light from the chair along the ray 92. However, in that case the mirror is a specular reflector, and so only one true reflected ray exists. That true reflected ray is illustrated as a broken line denoted by reference numeral 96 in Figure 5. In the example embodiment, no ray within the light field exactly coincides with true ray 96. Therefore, the parallel sub-field 70 in a direction nearest to the direction of the true ray 96 is identified and the ray 98 in that parallel

sub-field 70 closest to the true ray 96 is identified. Then the unshot radiance along the ray 92 is reflected along the best fit reflected ray 98. Therefore, by iterating through all objects in receipt of unshot radiance, radiance can be added at T-Intersections in accordance with the distribution of light in the scene. In fact, the objects marked as being in receipt of unshot radiance are considered in turn, in decreasing order of the amount of unshot radiance received by each object.

10 In that way, objects having most effect on illumination of the scene are dealt with first. The above technique is carried out without regard to any selected viewing position at this point.

15 The above examples as illustrated in Figures 5 to 9 have been described with regard to solid objects which reflect light either specularly or diffusely, or a combination of both. However, objects which are at least partially transparent can be treated in the same way. An object with transmissive properties will be described in terms

20 of the effect which the object has on light incident thereon. Radiance travelling along a ray intersecting a specularly transmitting surface of an object can be computed by identifying the point of exit of the light

25 from the object and the true direction thereof, taking account of, for example, the refractive properties of the object. The closest virtual light field ray to this true

direction can then be found and radiance can then be propagated along that closest ray.

5 The incidence of light on a diffusely transmissive object can be computed in the same way, but taking account of the fact that radiance will need to be propagated along a plurality of rays.

10 A virtual eye 44 is illustrated in Figure 10. The virtual eye 44 comprises a pupil 76, a lens 78, a focal plane 80, an image plane 82 and an optical axis. In order to view the scene from this point, all rays identified as entering the pupil 76 are considered to have entered the virtual eye 44. The direction of each ray and its position relative the optical axis 84 are identified, and a lens equation defining the structure and position of the lens 78 is applied thereto. A lens equation is a vector equation which identifies the trajectory of a refracted ray from a given trajectory of an incident ray.

15

20

Thereafter, the radiance associated with that ray at that point along the ray is recorded in an array associated with pixels or sub-pixels in order to build up an image.

25 Smoothing functions can be applied to the image so that any parts thereof which are not fully constructed having regard to the number of rays entering the pupil 76 can

be filled. It will be appreciated that the step of building up an image is carried out independently of the calculation of the characteristics of the light field with object and light emitters rendered therein, and so
5 that use of the light field is entirely view-independent.

The apparatus of the particular embodiment of the present invention will now be described in more detail with reference to Figures 11, 12 and 13.

10

Figure 11 is a basic block diagram illustrating operational modules of the computer 2 illustrated in Figure 1.

15

A user interface 100 comprises a facility for the interaction of an user with the apparatus. By virtue of the user interface 100, a user can identify the objects which are to be placed in a scene, and their characteristics such as light emission, absorption,

20

colour and/or position etc. Coupled with the user interface 100 is a light field (LF) computation unit 102.

The LF computation unit 102 is operative to define a light field in terms of the four-dimensional co-ordinate system defined above, to render objects into a scene

25

within the light field and to activate light emission such that radiance is distributed through the light field.

A viewer 104 is coupled with the LF computation unit 102 so as to obtain data therefrom concerning radiance values along rays 72 of the light field. The viewer 104 is operative to convert those radiance values into a focused
5 image, and to process that focused image into a form which can be viewed. The viewer 104 is linked with the VDU 12 as described with reference to Figure 1.

Figure 12 comprises a block diagram showing the
10 components of the LF computation 102. The LF computation comprises a preordained LF which cannot be modified by the action of the user. In an alternative embodiment, the resolution of the light field could be modified by the user. This could be carried out to select high
15 speed, or high quality imaging depending on user requirements.

Via the user interface 100, described with reference to Figure 11, a user defines, or calls up, a predetermined,
20 scene definition file 108. That scene definition file contains information concerning the nature of objects to be placed in the light field 106. Those object definitions are placed in an objects file 110.

25 The LF computation unit 102 further comprises an intersection data computation unit 112 which considers each object in turn from the objects file 110 in respect

of the preordained LF 106, to establish the position of intersections along each intersected ray of the LF. The intersection data computation unit 112 is operative to produce an intersection data table 114 for all rays
5 within the preordained LF 106.

The LF computation 102 also comprises a light computation unit 116 which makes reference to the objects file 110 to identify light emitting objects, and refers to the
10 preordained LF 106 and the intersection data table 114 to render light through the LF. Light is rendered through the LF, as described previously, by up-dating data within the intersection data table 114.

15 Figure 13 illustrates in more detail the components of the viewer 104. The viewer comprises a file of predetermined viewers 118 (such as a human eye, a regular camera lens, fish eye lens, wide-angle view, zoom etc) which can be selected and/or modified by interface with
20 the user through the user interface 100. The user interface 100 also provides a facility for specifying the position and orientation of the viewer 104 within the LF. The viewer 104 further comprises a viewed image computation unit 120 which makes reference to the
25 selected viewer lens and the intersection data table 114 of the LF computation 102. The viewed image computation unit 120 is operative to produce an image in accordance

with a lens equation describing the characteristic of the selected lens. It will be appreciated that a selected viewer 118 could include a sequence of lenses which, in combination, provide a desired optical effect.

5

Further, the viewer could comprise a light meter capable of delivering data representative of light received at a point, or over a surface. This could be particularly useful in architectural and theatrical design, where
10 light distribution is an important consideration.

Data corresponding to that image is passed to a data conversion component 122, such as for converting the data into raster scan data for use with a VDU. Up to that
15 point, the level of resolution of the image can be defined to suit the resolution of the light field, and need only be converted into a fully pixellated image at the final stage. The converted image is then output to the VDU 12.

20

A plurality of viewers 104 can be defined, either to deliver different images to different VDU's or, in the case of two viewers being provided, to deliver stereoscopic images to screens in a headset for use in
25 Virtual Reality applications.

Figures 14 to 20 describe procedures performed by the LF

computation 102 during operation thereof. Figure 14 describes the main procedure of the LF computation 102.

On commencement of the procedure, the LF is set up in
5 step S1-2 by the SET UP THE LF procedure described later. Once the LF is set up, objects are rendered into the LF in step S1-4 by the RENDER OBJECTS INTO LF procedure to be described later.

10 Once all objects have been rendered into the LF, those objects which are light emitters are activated in step S1-6, by means of the ACTIVATE LIGHT EMITTERS IN LF procedure to be described later. Then, in step S1-8, once the light emitters have been activated, radiance is
15 emanated through the LF by means of the COMPUTE REFLECTIONS IN LF procedure described below. Once that procedure has been completed, the LF is fully set up with objects rendered therein and light emission activated.

20 The SET UP LF procedure will now be described with reference to Figure 15. Firstly, in step S3-2, a light region 70 is defined as a grid of n by n parallel rays. The rays are arranged in x , y , z space parallel with the z axis and the light region is bounded by the following
25 constraints:

$$- 1 \leq x \leq 1$$

$$- 1 \leq y \leq 1$$

$$- 1 \leq z \leq 1$$

5 Following that step, in step S3-4 the light region is
 cycled through a predetermined number of orientations of
 ϕ and θ , wherein ϕ is the angle of a ray from the z axis
 and θ is the angle defined between a ray resolved into
 the x, y plane and the x axis. In order to cover all
 10 possible directions of the rays, ϕ and θ must pass
 through the following range:

$$\phi = 0$$

$$0 < \phi < \pi/2 ; 0 \leq \theta < 2\pi$$

15 $\phi = \pi/2 ; 0 \leq \theta < \pi$

ϕ , θ are discretised in accordance with the second method
 identified above, so as to generate a near uniform
 distribution of ray directions. In step S3-6, a set of
 20 equispaced rays parallel with the pre-rotated z-axis is
 generated, the rays being defined as look-ups into a
 four-dimensional array of interval trees as described
 above with reference to Figure 9. Step S3-8 enquires as
 to whether any further constructions of ϕ , θ need to be
 25 considered. If so, step S3-10 selects the next
 combination, and the cycle recommences. If not, the
 routine returns.

The RENDER OBJECTS INTO LF procedure enquires in step S5-2 whether any objects are to be rendered into the LF. This is carried out by calling the OBJECTS file 110. If there are no objects to be rendered into the LF, the procedure returns to the main procedure. Otherwise, in
5 step S5-4, the OBJECTS file 110 is called to obtain the parameters defining the object. Those parameters consist of the shape definition, which may be in terms of its geometry, or a skeleton, and the surface characteristics
10 of the object. Those surface characteristics may consist of the light absorption characteristics, and the bi-directional reflectance distribution function (BRDF). The object parameters may also include details of any light emission characteristics of the object. At this
15 point, however, only the geometry of the object is taken into account.

Following the obtaining of the parameters defining the object, the procedure calls, in step S5-7, a sub-
20 procedure namely OBTAIN DETAILS OF INTERSECTS OF RAYS WITH OBJECTS 210. When that sub-procedure has completed, the procedure returns to step S5-2 to repeat the enquiry as to whether any objects remain to be rendered into the LF. If no more objects require rendering into the LF,
25 then the procedure returns, otherwise the procedure selects the next object and repeats as necessary.

The OBTAIN DETAILS OF INTERSECTS OF RAYS WITH OBJECT procedure will now be described with reference to Figure 20. Firstly, in step S13-2, the procedure enquires as to whether any rays intersecting the object remain to be considered. If no more rays remain to be considered, the procedure returns to the RENDER OBJECTS INTO LF procedure described above. Otherwise, in step S13-4, an intersected ray is considered and, in step S13-6, intersections of that ray with the object are successively loaded onto a binary tree for that ray. The intersection data consists of the position of the intersection on the ray, the identity of the object intersected, the potential direction of radiance from the intersection, and radiance and unshot radiance values which are initially set to zero. The procedure then enquires in step S13-8 as to whether any other rays remain to be considered. If no more rays remain to be considered, then the procedure returns to the earlier described RENDER OBJECTS INTO LF procedure. Otherwise, the procedure loops back to step S13-4.

Once all objects have been rendered into the LF, the ACTIVATE LIGHT IMAGES IN LF procedure is called. This procedure is described in more detail with reference to Figure 17. The procedure commences in step S7-2 by considering an object with light emission properties. Thereafter, a sub-procedure FIND A RAY INTERSECTING

OBJECT AND INTERSECTION THEREWITH 212 is called in step S7-4. This procedure will be described later.

Once a ray intersecting the object and its intersection
5 therewith has been identified, in step S7-6, an interval
is defined which is bounded by that intersection and the
next intersection along the ray from that object in the
direction of predicted radiance (i.e. "right" or "left")
defined in the data associated with the intersection.
10 Then, in step S7-8, the radiance in that interval is set
in the data structure associated with the intersection
under consideration. That radiance is set in accordance
with the light emitting properties of the object, which
are called from the parameters defining the OBJECT as
15 held in the objects file 110.

In step S7-10, an enquiry is made as to whether radiance
in that interval is greater than a predetermined
threshold. If so, then the procedure concludes that it
20 would be worthwhile computing any reflections from that
radiance. As such, step S7-12 then causes the unshot
radiance in the intersection data in question to be set
so as to correspond with the change in the radiance. The
object at the other end of that interval is identified
25 in step S7-14 and that object is marked in step S7-16 as
being in receipt of unshot radiance.

If the radiance in the interval is less than the threshold, then the unshot radiance is not set, and there is no requirement to identify and mark the object at the other end of the interval. Thereafter, any other
5 intersections of that ray with the object in question are considered in steps S7-18 and S7-20 and radiance and unshot radiance are updated as required. Any other rays intersecting the object are then considered in step S7-22, and then in step S7-24 any other objects with light
10 emission properties are considered in the same way. Once all objects having light emission properties have been processed in the same way, the procedure returns.

The ACTIVATE LIGHT EMITTERS IN LF procedure called the
15 FIND RAY INTERSECTING OBJECT AND INTERSECTION THEREWITH procedure. That procedure will now be described with respect to Figure 19. The procedure firstly enquires in step S11-2 as to whether the object is defined in terms of plane polygons. That information is contained in the
20 OBJECTS file 118. If it is so described, a more simplified procedure is followed. A plane polygon is considered in step S11-4, and a ray passing through the vertices of the polygon is found in step S11-6. The ray is marked in step S11-8 as having been found, and the
25 intersection of that ray is found in step S11-10 by Z-interpolation within the polygon under consideration. Once that intersection is found, the procedure returns.

The procedure of finding and dealing with rays passing through a plane polygon can be dealt with using a 2-D fill algorithm as set out in "Computer Graphics, Principles and Practice" by Foley, Van Dam, Feiner and
5 Hughes, pp 92-99. Clearly, the use of values of z , to define positions of intersection of rays with objects, also allows use of the invention alongside z -buffer techniques as also discussed in that publication, pp 668-672, to accommodate less photo-realistic but dynamic
10 objects within an illustrated 2D image of the scene.

In an alternative case, step S11-2 may identify that an object is described not in terms of plane polygons, but in terms of a geometric shape such as a sphere or an
15 ellipsoid, or a composite of various primitives.

One method of rendering such an object into the scene involves the definition of a bounding box around the object in step S11-12, the bounding box being axis
20 aligned and of smallest size possible while containing the object. A side of the box is considered in step S11-14, and a ray passing between the vertices of that side of the box is found in step S11-16. As each ray is found, a check is made in step S11-18 as to whether the
25 ray has already been marked as found. If so, an enquiry is made in step S11-20 as to whether any more rays need to be considered. If not, a new side of the box is

considered in step S11-14. If so, a new ray is found from the same side in step S11-16. If the enquiry of step S11-18 shows that the ray is not marked as found, then in step S11-24 an enquiry is made as to whether the ray passes through the object. If it does not pass through the object, then the procedure continues from step S11-20, and the same checks are carried out as before.

10 Once a ray has been found, which passes through the object, then the intersections of the ray with the object are calculated with respect to the object shape definition. Once those intersections has been identified, then the procedure returns.

15

The intersections between rays and the object can be identified in an alternative manner. For every combination of ϕ and θ , a 2D projection of a bounding box around the object is found. Then, all rays passing through that 2D projection in the direction of ϕ , θ are tested against the object for intersection (if any) and the point along the ray at which intersection takes place.

20

25 Finally, once all light emitters in the LF have been activated, and radiance has been added to all of the rays intersecting those light emitters, reflections through

the LF are computed. That is carried out by calling the COMPUTE REFLECTIONS IN LF procedure described now in more detail with reference to Figure 18. Firstly, in step S9-2, the procedure inquires as to whether there are any
5 objects marked as having received unshot radiance. Objects will have been marked as a result of ray intersections having had radiance added to them above a predetermined threshold, and therefore corresponding unshot radiance will have been added thereto as well.
10 If no such objects are found, the procedure returns.

An object which has been so marked is put under consideration in step S9-4, and is then unmarked in step S9-6. In step S9-8, the procedure FIND RAY INTERSECTING
15 OBJECT AND INTERSECTION THEREWITH is called. The intersection T_2 of that ray with the object O is then identified in step S9-10. The procedure then searches in step S9-12 through the data structure, which is conveniently in binary tree form, to identify the
20 adjacent intersection T_1 along the ray in the direction of probable radiance as defined in the intersection data for intersection T_2 . The adjacent intersection T_1 is checked in step S9-14 to establish whether it contains unshot radiance. If it does not contain unshot radiance,
25 then the procedure returns to step S9-8 to consider another ray intersecting the object in question as

before. Once a ray has been identified which has an intersection with the object and which is an adjacent intersection containing unshot radiance, a reflected ray is identified.

5

The step of finding a reflected ray in step S9-16 requires reference to the BRDF of the object, which is contained in the OBJECTS file 110. The BRDF may specify diffuse reflectance, specular reflectance or a combination of both. In the case of diffuse reflectance, a plurality of reflected rays exist to be found, over the point of intersection of the incident ray under consideration with the object. With a specular reflection, the reflected ray may be a ray of best fit having regard to a true reflected ray as calculated with respect to the well-known laws of reflection.

In addition to specular and diffuse reflection, the BRDF may specify more complicated reflection models such as "glossy" reflection, where a cone of reflected rays are produced to correspond with a given incident ray, this effect represents imperfect specular reflection, as described by Phong in CACM 18(6), June 1975, 311-317.

In respect of the or each reflected ray, the intersections of that ray with the object in question are identified in step S9-18. In step S9-20 the change in

radiance in accordance with the BRDF is considered as to whether it is greater than the aforesaid predetermined threshold. If it is greater than the predetermined threshold, then, in step S9-22, the radiance in accordance with the BRDF is set in the intersection with the object in question. Then, the aforesaid adjacent object is marked in step S9-24 as having received unshot radiance, and the aforesaid intersection with the object under consideration is updated in step S9-26 so as to reflect the increase in unshot radiance. Otherwise, no consideration is taken of unshot radiance.

An enquiry in step S9-28 is then made as to whether any more reflected rays remain to be considered. If so, the procedure returns to step S9-10. Once that has been carried out in respect of all reflected rays to be considered, the object is further considered in step S9-30 to establish whether any other rays intersecting that object have received unshot radiance. If so, then the procedure returns to step S9-8. If no more rays are to be considered, then the procedure returns to step S9-2 where the light field is further considered to establish whether any objects remain which are marked as having received unshot radiance. Once all objects are unmarked, the procedure returns.

Once all unshot radiance from all objects has been

considered, the light field is fully defined with respect to the scene rendered therein. At that point, a complete data structure exists independent of any view position or the characteristics of any viewer. It would be possible to use the light field without reference to any viewer, for instance in the design of an art exhibition where even distribution of light throughout a region of the light field needs to be considered. The distribution of light in a gallery could be designed with the above-described arrangement, and the light incident on any particular wall could be monitored. In that way, bright patches or dark patches could be eliminated from a particular wall.

However the present invention is also particularly suitable for applications where a particular viewpoint is required, such as is illustrated in Figure 10 with the simulated eye 44. A view can be constructed from the LF through the procedure now to be described with reference to Figure 21. This procedure stands alone from the procedure operating the LF computation, and can be controlled by the user interface 100. The viewer procedure commences by calling for a specification of an aperture of the viewer. This can be identified from the VIEWERS file 118 under the control of the user 100. The specification of the aperture includes the actual diameter of the aperture plus its position in space and

orientation. Once the aperture has been specified, a set of rays entering the aperture can be identified. This can be carried out using a procedure such as the FIND RAY INTERSECTING OBJECT AND INTERSECTION THEREWITH procedure previously described.

As each ray is found which enters the aperture, the end point of that ray on the image plane can be found. The ray can be considered as a vector which can be passed through a vector equation relating to the composition of the lens. The composition of the lens can be described with reference to the VIEWERS file 118. The result of the vector equation is a point on an image plane 82. The image plane is defined as an array of elements, and the radiance associated with the segment of that ray, as called back from the previous intersection of the ray with an object, is mapped onto an element within that array. It is possible that several rays will be incident on any particular element, and as each ray is incident on an element, the radiances are combined together.

Eventually, as all rays are considered, an intensity map will be formed on the image plane composed of elements, which can be processed by filtering and transformation from the array of pixels to form a pixellated image. In particular, if the lens structure is that of a relatively simple box camera, the image will need inversion before

it can be displayed correctly on a VDU.

The viewing procedure set up above can be relatively fast, and is independent of the level of complexity of the image. In fact, it can be considered to be generated in substantially constant time from frame to frame as the viewer moves. With sufficient and reasonable processing power, that constant time could be reduced to such an extent that the viewer could be implemented as a real-time viewer.

The above identified arrangement can also be used in a headset for virtual reality applications, with stereoscopic images being projected to the eyes. Both images can be computed rapidly. Moreover, the orientation of the viewing direction from which images are computed can be altered depending on the angle of gaze of the eyes of the wearer of the headset. This is an important development which has previously required significant computational power for its achievement.

Although the above apparatus has the potential for a view of a static scene to be changed in real-time, dynamic changes to the scene itself in real-time may be rather more difficult to achieve. However, the apparatus can be combined with other apparatus to define an image of a static, background scene in accordance with the present

disclosure, while the other apparatus produces an image which can be changed dynamically, which image is superimposed over the image of the static scene. As noted earlier, the apparatus can readily be used to
5 retrieve z-values for combination with a z-buffer method for rendering objects into a scene.

Moreover, the embodiment described above assumes that the light transmitting medium is non-participatory. In
10 practice, media are rarely non-participatory, and so attenuating properties can be represented by means of placing intersections at random along the rays in the light field. The concentration of those random intersections depends on the level to which the medium
15 in question attenuates light transmitted through it. In that way, the scene can be illuminated taking account of the attenuating properties of the medium. In fact, attenuation is commonly the result of particles (such as water droplets or dust) in suspension in air, and so the
20 random distribution of instructions is a reasonable approximation of reality.

The preferred embodiment of the invention as described above has been described in relation to the modelling of
25 a scene to which no changes are to be made following modelling. However, circumstances could arise in which it would be desirable to be able to delete an object from

the scene, or to add an object to the scene. Figures 22 and 23 describe procedures by which these two actions can be achieved.

5 Firstly, Figure 22 illustrates a flowchart defining a procedure for deleting an object from a scene defined in the light field to which light has already been applied. In step S17-2, a ray is identified which intersects the object to be deleted. Thereafter, in step S17-4, a
10 T-Intersection is found for the identified ray and carrying radiance incident on the object in question.

Then, in step S17-6, the radiance R along that interval is noted. The surface normal at that second intersection
15 on the object to be deleted is found, and the redistribution of radiance along reflected rays is determined by adding negative radiance (and unshot radiance) to T-intersections for the or each reflected ray. Unshot radiance should be dealt with in accordance
20 with previously described techniques, especially with reference to Figure 18.

Then, in step S17-8, the radiance at the T-intersection corresponding with the transmitted ray is set to equal
25 the incident radiance R .

Then, an enquiry is made in step S17-10 as to whether any

more T-intersections are to be considered in respect of the object to be deleted. If there are, then the procedure repeats in respect of the next T-intersection.

- 5 Then, an enquiry is made in step S17-12 as to whether any further rays are to be considered. If so, the above procedure is repeated in respect of further rays; otherwise, in step S17-16, the intersections corresponding to the deleted object are deleted from the
- 10 interval trees relating to those rays. Then the previously described routine COMPUTE REFLECTIONS IN LF can be called, in relation to objects marked as being in receipt of unshot energy.
- 15 Figure 23 is a flow diagram showing steps of the procedure designed to allow an object to be added to a scene rendered into a light field, after light emitters have been activated.
- 20 Firstly, in step S19-2, a ray intersecting the object to be added is found. Then, in step S19-4, an interval along the ray is identified, bounded by a so-called second intersection with the object to be added, the direction of potential radiance from that second
- 25 intersection being to the left. The radiance R currently along that interval is noted in step S19-6, and the surface normal at the second intersection with the object

is found in step S19-8. Rays along which that radiance R is to be propagated are determined in step S19-10, and in step S19-12 radiance and unshot radiance are added to the intersections of those rays with the object as
5 necessary. Any objects in receipt of unshot energy are marked in step S19-14.

The interval along the ray corresponding to the inside of the object is identified in step S19-16, and in step
10 S19-18 the radiance therealong is set to zero.

Finally, an interval along the ray is identified in step S19-20, bounded by a so-called first intersection with the object, the direction of potential radiance from that
15 first intersection being to the right. The unshot radiance along that interval is set to be $-R$ in step S19-22, and the object (if any) at the other end of that interval is marked in step S19-24 to be in receipt of unshot energy. The fact that the unshot energy received
20 by the object is "negative" is not relevant; it is only important that there is a change in unshot energy which needs to be processed through the light field.

Then, an enquiry is made in step S19-26 as to whether any
25 further rays are to be considered in relation to the object to be added, and if so, the above steps are repeated. If no rays remain to be considered, the

procedure COMPUTE REFLECTIONS IN LF is called in step S19-28 in respect of the objects marked as being in receipt of unshot energy (whether positive or negative), following which the ADD AN OBJECT procedure is completed.

5

The above two procedures are optional features, which can preferably be called from the user interface 100.

These two procedures can be combined in order to provide
10 a procedure for moving an object. In the combined procedure, the union of the object in its original position and the object in its new position is put under consideration, and the rays intersecting this union are found. This is computationally less expensive than
15 considering the deletion and the addition of the object as separate steps to be applied in turn.

The invention, as described by way of the above exemplary apparatus and procedure, is particularly advantageous in
20 that it is capable of providing a globally illuminated scene which can be amended by adding, deleting or moving an object in the scene without the need to recalculate the illumination of the entire scene.

25 Further and alternative means of listing the T-Intersections could be provided. For example, an array could be defined for each ray and the T-Intersections

could be loaded into this array by t-value. This has the disadvantage of needing to choose a maximum number of intersections along a ray in advance.

- 5 Alternatively, the T-Intersections could be stored in a standard linked list, or a doubly-linked list so that it is particularly easy to insert new elements or to delete elements from the list. The disadvantage of this is that there would be linear search time for any particular
10 element, compared to the logarithmic search time required for the binary tree method.

- As a final alternative, it would be possible to divide the ray into N equal segments. If N is large (say 1000)
15 then any t value is approximated by an index into a particular segment of the ray. Hence, a one dimensional array of T-Intersections, indexed by the index number of the segment corresponding to the intersection, could be used. This has the advantage of constant look-up time
20 (so it is faster than the binary tree method for large number of intersections) but the disadvantage of approximation and memory requirements (since most of the entries in the array would be null). For example, suppose that $N = 1000$, but that there are only twenty T-
25 Intersections along a given ray. In that case, 980 of the possible entries in the array for that ray would be empty.

Two ways exist of overcoming the memory disadvantages of the final alternative set out above. The first way is to use a run length encoding method. In that case, the storage would be in the form of the number of null
5 entries, followed by the actual non-null entries, and so on throughout the array. However, the look-up time for that array would be non-constant depending on the number and location of non-null entries of the array.

10 The second way of improving the final alternative set out above is by packing many possible representations into words, rather than considering them as single array entries. For example, if $N = 1024 (2^{32})$, 1024 entries of "0" and "1" can be represented by 32 unsigned integers.
15 Hence, instead of having an array of 1024 potential T-Intersections, an array of 32 unsigned integers, each initialised to zero, is provided. This is less memory intensive. Associated with each such integer is an ordered list of those particular T-Intersections that
20 correspond to "1" entries in the word. So, given a particular t-value, the closest segment in the representation can be found. The quotient of the index of the segment divided by 32 will give a particular array element in which the t-value is located. The remainder
25 of the index of the segment divided by 32 will identify the particular bit that must be set to 1 to represent this t-value. Finally, the T-Intersection can be stored

in a linked list associated with this particular integer.

It will be appreciated that the techniques described above are computationally relatively more expensive when considering diffuse reflectors than when considering specular reflectors. That is because diffuse reflection produces a very large number of reflected rays from a single incident ray. If a diffuse reflector is included in a scene and dealt with according to the method described above, the level of computation can become prohibitive. Therefore, the following procedure has been devised with a view to significantly reducing the level of computation required for considering diffuse reflectors within a scene.

For example, Figure 24 shows a perfectly diffusely reflecting object 230, a surface of which is intersected by first and second rays 232, 234 of a light field as described above. The two rays 232, 234 also intersect another body (not shown) which is a light emitter. The two rays 232, 234 intersect a surface of the diffusely reflecting object 230 substantially at the same point.

Therefore, when considering the receipt of unshot energy by the object, and consequent reflection of light, according to the previously described method the first ray 232 is considered and its intersection with the

diffuse reflector 230 is identified. Once the intersection has been found, reflected rays are identified through the hemisphere centred at the intersection and defined by the region bounded by the plane tangential to the surface of the object at the intersection. Radiance and unshot radiance are applied to those reflected rays. Thereafter, the second ray 234 is considered in the same way. However, since the intersection of the second ray 234 with the diffuse reflector is substantially the same as the intersection of the first ray 232 with the diffuse reflector, many, if not all, of the reflected rays in respect of the second incident ray will be the same as those for the first incident ray.

15

The following method has been devised with a view to reducing the amount of computational effort required in transmitting radiance resultant at a diffusely reflective surface.

20

Figure 25 shows a schematic diagram of a portion of the surface of the diffuse reflector 230. By way of explanation of the method, the portion of the surface has been divided into a grid comprising a plurality of elements 240. The two incident rays 232, 234 are illustrated as intersecting the object surface within one particular element 240. Once the first ray intersection

has been found, instead of reflecting the ray immediately, the energy associated with that incident ray is accumulated with respect to that element 240. Then, when the second incident ray 234 is calculated as
5 intersecting within the same element, the energy associated with that ray is also accumulated in respect of that element 240. This step results in elements of the surface having accumulated energy associated therewith. The accumulation of energy is carried out
10 according to a weighted sum of terms. Each term represents energy received from a different ray. The weighting of each term relates to the angle of incidence of the incident ray on the object surface, and the properties of the surface as defined by the BRDF for the
15 surface.

Figure 26 illustrates the same portion of the object surface, in which a step is illustrated which sends radiance onto rays identified as being in reflection
20 directions from the surface element 240. In this step, the accumulated energy for the surface element 240 is reflected to all identified reflected rays from the surface element 240. In that way, the computationally expensive reflected energy calculation is carried out
25 only once per surface element 240. Therefore, the level of computation required for calculating diffuse reflection can be reduced.

Accumulated energy for an element 240 is stored in an additional field within the structure of the T-intersection associated with that element and relating to a ray of the light field substantially coincident with the surface normal of the object surface at that element of the surface. Practically, it is not possible to add the field only to the set of T-intersections corresponding to surface normals, and so the field is included as part of the structure of T-intersections in general. For all T-intersections the accumulated energy field is initialised to zero. For all but the T-intersections corresponding with surface normals, the accumulated energy field will remain set to zero. Therefore, as shown in Figure 27, the T-intersection associated with the ray coincident with the surface normal N at the point of intersection of the incident rays 232, 234 has stored therein the accumulated energy related to the radiance of the incident ray.

Thereafter, the accumulated energy is released onto the reflected rays calculated in accordance with the BRDF for the object at that point.

By storing information relating to accumulated energy per element 240 in a part of the T-intersection data structure, one particular advantage of the specific embodiment of the invention is maintained. That is,

there is no data held with respect to objects in the scene, only in terms of intersections.

The present invention has several applications, some of which have been identified throughout the description. Other applications will now be described.

Firstly, a model of a building can be rendered into the light field, and illuminated in a desired manner, to verify the quality of an architectural design. Moreover, lighting can be designed within that rendered building, for use in the production of artistic works, such as exhibitions, stage plays, or film sets.

Secondly, the ability to interchange views and view characteristics enables the use of the invention in the design of a sequence of shots for a television or film production. The invention would be suitable to simulate in advance of constructing a film set the appearance of that set under various layout and lighting conditions, and with different types of camera lens.

Thirdly, the light field is a digitally encoded representation of a scene, through representation of the light within the scene. A digital encoding is ideal for compression and subsequent transmission, for instance, for broadcast television. A local set top box, hardwired

to decode the transmitted digital encoding and capable of extracting images from the light field defined thereby could be used to display games, or other entertainment to people in their homes. The digital encoding could
5 also be transmitted across the Internet.

Fourthly, in the case that two images are provided of views of the same light field, those two images can be transmitted to a head mounted display. In that way,
10 stereoscopic images can be formed, which can provide a highly immersive virtual reality environment.

In Figure 28, a block diagram is shown of an image processing apparatus having all of the components of the
15 image processing apparatus illustrated in Figure 1.

However, in this case, the central processing unit 4 is operative to produce two images of the light field embodied therein. Those two images are produced by two
20 viewers within the light field, having parallel principal optical axes a suitable interpupillary distance apart, which interpupillary distance can be adjusted by a user.

The images produced at those viewers are delivered to the
25 frame buffer 10 from the central processing unit 4. Coupled with the frame buffer 10 is a CD manufacturing suite 250 operative to produce compact discs 252 on which

are stored encoded data relating to images delivered from the frame buffer 10. In that way, sequences of images from a light field in accordance with the invention can be stored on compact discs for viewing later on suitable viewing apparatus.

A head mounted display (HMD) 254 is provided, and is operative to receive both sequences of images from the frame buffer 10. A second VDU 12' is provided, and the two sequences of images can be displayed on those two VDU's 12, 12'.

In this case, the input means 8 can include a manual input device, such as is used commonly in virtual reality applications. In that way, progress of the user through the virtual reality environment can be controlled by the user through manipulation of the manual input device, and the two images produced by the central processing unit 4 can be adjusted accordingly. Moreover, the interpupillary distance can be measured within the HMD 254, and that measurement can be supplied to the input 8. That interpupillary distance can be translated into the distance between the optical axis of the two viewers within the light field defined in the central processing unit 4 which produce the two images for display in the HMD 254.

The CPU 4 is further coupled to a data compressor 260 which is adapted to receive data relating to the definition of the light field in the CPU 4 and compressing according to a suitable compression
5 algorithm. The data compressor 260 is coupled to a transmitter 262 operable to receive compressed data from the data compressor 260, and to transmit that data by modulation on electromagnetic carrier radiation. The transmitted data is received by a receiver 264 which is
10 in practice implemented as a set top box for a television 266. The receiver 264 comprises a decoder 268, capable of receiving transmitted data and returning it to its uncompressed state, to retrieve the data defining the light field. The data defining the light field is then
15 transferred to a viewer 270 of the receiver 264 which is of a similar construction to the viewer 104 illustrated in Figure 13 of the drawings. In that way, a definition of a 3D environment can be transmitted to a receiver for manipulation by a user.

20

With reference to Figure 29, the HMD 254 consists of a visor suitable to be placed over a user's head 272. A cable 274 connects the HMD 254 to the image processing apparatus previously described.

25

The HMD 254 conventionally comprises two display units 276 (as illustrated in Figure 30) on which the previously

described stereoscopic images can be projected.

The present invention is particularly suitable for viewing an environment in stereo, since there is little computational effort involved in creating an image from a light field which has already been constructed. Therefore, two views of the same light field can be created without excessive computation. The viewers can be altered in response to monitoring the angle of gaze of the eye of the user, using monitoring means incorporated into the head mounted display 254, to establish objects within the scene which are of interest to the user. In that way, those objects can be brought into focus.

Furthermore, the invention is not limited to the representation of light propagation in a scene. Heat and sound also constitute forms of energy which can be represented as radiating from sources, and so the representation of the propagation of those forms of energy could also be accomplished by the above described technique. In fact, additional fields could be included for heat and/or sound data, in the T-intersection data structure, beyond the existing radiance and unshot radiance fields. In accordance with the invention, heat and unshot heat fields could be added, as could sound and unshot sound fields.

The invention is suitable for dynamic sources of light, heat and/or sound, since the radiance value at a particular T-intersection can be computed as a function of a nominal source intensity, until the actual source
5 intensity is set. When the actual intensity is set, the actual radiance value at that T-intersection can be ascertained. Therefore, in the case of a field being used for sound propagation determination, propagation may be calculated as a function of a nominal sound intensity,
10 and a received sound intensity can be calculated thereafter by computing the output of that function given a desired sound source intensity.

CLAIMS:

1. Apparatus for analysing propagation of radiative energy in an environment, the apparatus comprising:

5 first storage means for storing information defining a plurality of paths in a plurality of directions in the environment;

second storage means for storing information relating to energy propagating along said paths in said scene, wherein said second storage means is operative to
10 store first and second information, said first information relating to intersection of at least one path with an object within said environment, and said second information relating to propagation of energy at said
15 intersection.

2. Apparatus in accordance with claim 1 for graphically representing a scene in an environment, the second storage means being capable of storing information
20 relating to light energy propagation along said paths.

3. Apparatus in accordance with claim 2, and further comprising processing means for identifying a source of light in said scene and for processing information
25 relating to illumination of light from said source of light, said processing means being operative to act on said information relating to illumination of light and

said first information to generate said second information taking account of the existence of said light source.

5 4. Apparatus in accordance with claim 1 or claim 2,
 wherein said processing means is operative to generate
 further information relating to intersection of at least
 one of said paths with a further object to be introduced
 into said environment, and to process said first
10 information to generate new second information taking
 account of the introduction of said object.

 5. Apparatus in accordance with any preceding claim,
 wherein said processing means is operative to remove from
15 said first information relating to intersection of said
 path or paths by an object to be removed from
 representation in said environment, removing said first
 information relating to that object from said second
 storage means, and generating further second information
20 relating to propagation of energy along paths taking
 account of the removal of said object.

 6. Apparatus in accordance with any preceding claim,
 wherein said first information defines paths such that
25 the distribution of path directions within the
 environment is substantially uniform.

7. Apparatus in accordance with claim 6, wherein the first information defines a direction of a path relative a reference plane in the environment, such that a larger number of paths are defined in directions at smaller angles to the reference plane than are defined in directions at larger angles to the reference plane.

8. Apparatus in accordance with claim 7, wherein the first information defines paths such that the number of paths defined in directions at a particular angle to the reference plane is substantially proportional to the complement of said particular angle.

9. Apparatus in accordance with any preceding claim, wherein the first information defines a plurality of paths comprising a plurality of sub-sets of paths, each sub-set comprising a sub-plurality of parallel paths.

10. Apparatus in accordance with claim 9, wherein said first storage means comprises means for indexing said sub-sets, said sub-sets being indexed in accordance with the direction of the paths contained therein.

11. Apparatus in accordance with claim 10, wherein each sub-set are indexed in accordance with spherical coordinates describing said direction.

12. Apparatus in accordance with any preceding claim, wherein said second storage means is operative to store, alongside said information relating to the intersections, information relating to the identity of the object in the scene with which a path intersects at that intersection.

13. Apparatus in accordance with any preceding claim, including image forming means for producing an image, said image being indicative of energy propagating along paths within the environment.

14. Apparatus in accordance with claim 13, wherein the image forming means is operative to position an aperture within the environment, and to determine an image indicative of energy propagated along paths passing through said aperture.

15. A storage medium comprising a set of processor implementable instructions for configuring a computer to operate as apparatus in accordance with any one of claims 1 to 14.

16. A signal carrying a set of processor implementable instructions for configuring a computer to operate as apparatus in accordance with any one of claims 1 to 14.

17. A method of analysing propagation of radiative

energy in an environment, the method comprising the steps of:

defining an environment within which energy can be propagated;

5 defining a plurality of paths in the environment along which energy can be represented as propagating within the environment; and

10 recording one or more points on a path, at which energy is to be represented as interacting with the environment under analysis.

18. A method in accordance with claim 17, for representing graphically propagation of light further comprising the steps of:

15 defining a scene in the environment, for graphical representation; and

storing information, in said recording step, relating to light energy propagating along said path at said point or points.

20

19. A method in accordance with claim 17 or claim 18, wherein said step of defining a plurality of paths includes the step of defining paths such that the distribution of path directions is substantially uniform.

25

20. A method in accordance with any one of claims 17 to 19, wherein the step of defining a plurality of paths

including defining a direction of a path relative a reference plane in the field, such that a larger number of paths are defined in directions at smaller angles to the reference plane than are defined in directions at
5 larger angles to the reference plane.

21. A method in accordance with any one of claims 17 to 20, wherein the step of defining a plurality of paths includes defining paths such that the number of paths
10 defined in directions at a particular angle to the reference plane is substantially proportional to the complement of said particular angle.

22. A method in accordance with any one of claims 17 to
15 21, wherein the step of defining a plurality of paths includes defining a plurality of paths comprising a plurality of sub-sets of paths, the paths within a sub-set being substantially parallel.

20 23. A method in accordance with claim 22, further comprising the step of indexing said sub-sets in accordance with the direction of the paths contained therein.

25 24. A method in accordance with claim 23, including the step of indexing the sub-sets in accordance with spherical coordinates.

25. A method in accordance with any one of claims 17 to 24, wherein said step of recording a point on one of said paths includes the step of recording the position of said point on said path.

5

26. A method in accordance with any one of claims 17 to 25, wherein said step of recording a point on one of said paths includes the step of recording a direction of potential propagation of energy along said path.

10

27. A method in accordance with any one of claims 17 to 26, wherein said step of recording a point on one of said paths includes the step of recording data representative of intensity of energy propagated along said one of said paths from said point.

15

28. A method in accordance with any one of claims 17 to 27, wherein said step of recording a point on one of said paths includes the step of recording the identity of an object in the environment with which the path interacts at that point.

20

29. A method in accordance with any one of claims 17 to 28, including the step of producing an image, said image being indicative of energy propagating within the environment.

25

30. A method in accordance with claim 29, wherein the step of producing an image includes the steps of:

defining an aperture within the environment; and

producing an image indicative of energy propagated
5 along paths passing through said aperture.

31. A storage medium comprising a set of processor implementable instructions for configuring a computer to operate in accordance with the method of any one of
10 claims 17 to 30.

32. A signal carrying a set of processor implementable instructions for configuring a computer to operate in accordance with the method of any one of claims 17 to 30.
15

33. A method in accordance with any one of claims 17 to 30, further including the step of providing a removable storage medium and recording on said storage medium information defining said one or more points on said one
20 or more paths, thereby defining interactions in said scene to be represented.

34. A method in accordance with claim 33, wherein said information includes information relating to light
25 propagation along portions of said paths between said points of interactions.

35. A method in accordance with any one of claims 17 to 30, and further including the step of transmitting information on a data carrier carrying a set of data defining information relating to said interaction within
5 said scene.

36. A method in accordance with claim 35, wherein said records include records of energy propagating along portions of said paths between said points of
10 interaction.

37. A method of transmitting a signal comprising the steps of:
generating a signal in accordance with claim 35 or
15 claim 36; and
transmitting said signal.

38. A method of constructing a graphical image, the method comprising:
20 receiving a storage medium produced in accordance with the method of claim 33;
defining a viewer aperture within the plurality of paths;
defining a viewer image plane in said plurality of
25 paths;
constructing an image defined by points of intersection of paths passing through the aperture with

the image plane.

39. A method of constructing a graphical image, the method comprising:

5 receiving a signal transmitted in accordance with claim 37;

defining a viewer aperture within the plurality of paths;

10 defining a viewer image plane in said plurality of paths;

constructing an image defined by points of intersection of paths passing through the aperture with the image plane.

15 40. A virtual reality system comprising:

apparatus in accordance with any one of claims 1 to 16; and

20 left and right image generation means, each image generation means being operable to produce an image from a viewpoint within the environment of energy propagated along paths towards the viewpoint.

41. A virtual reality system in accordance with claim 40, further including left and right display means for
25 displaying said images for stereoscopic viewing thereof.

42. Apparatus for monitoring energy distribution through

an environment comprising apparatus in accordance with any one of claims 1 to 16 and energy measuring means operable to measure energy propagating through a region of the environment.

5

43. Apparatus in accordance with claim 42, wherein the energy represented as propagating through the environment represents light, the apparatus being operative to measure light distribution in an environment.

10

ABSTRACTENERGY PROPAGATION MODELLING

A simulation apparatus (2, 70) is operable to simulate
5 propagation of light within a scene. The apparatus
comprises means for defining a plurality of discrete
paths (72) along which energy propagation in the
environment is to be traced, and means for recording
points (t) on the paths, at each of which points there
10 is an interaction within the scene with the path (72),
A corresponding method is also provided.


The invention further provides apparatus (2) for
monitoring light distribution in an environment, and
15 apparatus (2, 254) for displaying a stereoscopic image,
for instance for use in virtual reality applications.

Distribution of propagation of either energy, such as
sound or heat, can be simulated in addition to or as an
20 alternative to light propagation simulation.


(Figure 5)

This Page Blank (uspto)


FIG. 1



This Page Blank (uspto)



This Page Blank (uspto)



This Page Blank (uspto)

FIG. 3A




FIG. 3B






FIG. 3C




This Page Blank (uspto)



This Page Blank (uspto)

FIG. 5



This Page Blank (uspto)




FIG. 6

This Page Blank (uspto)




FIG. 7

This Page Blank (uspto)




FIG. 8

This Page Blank (uspto)

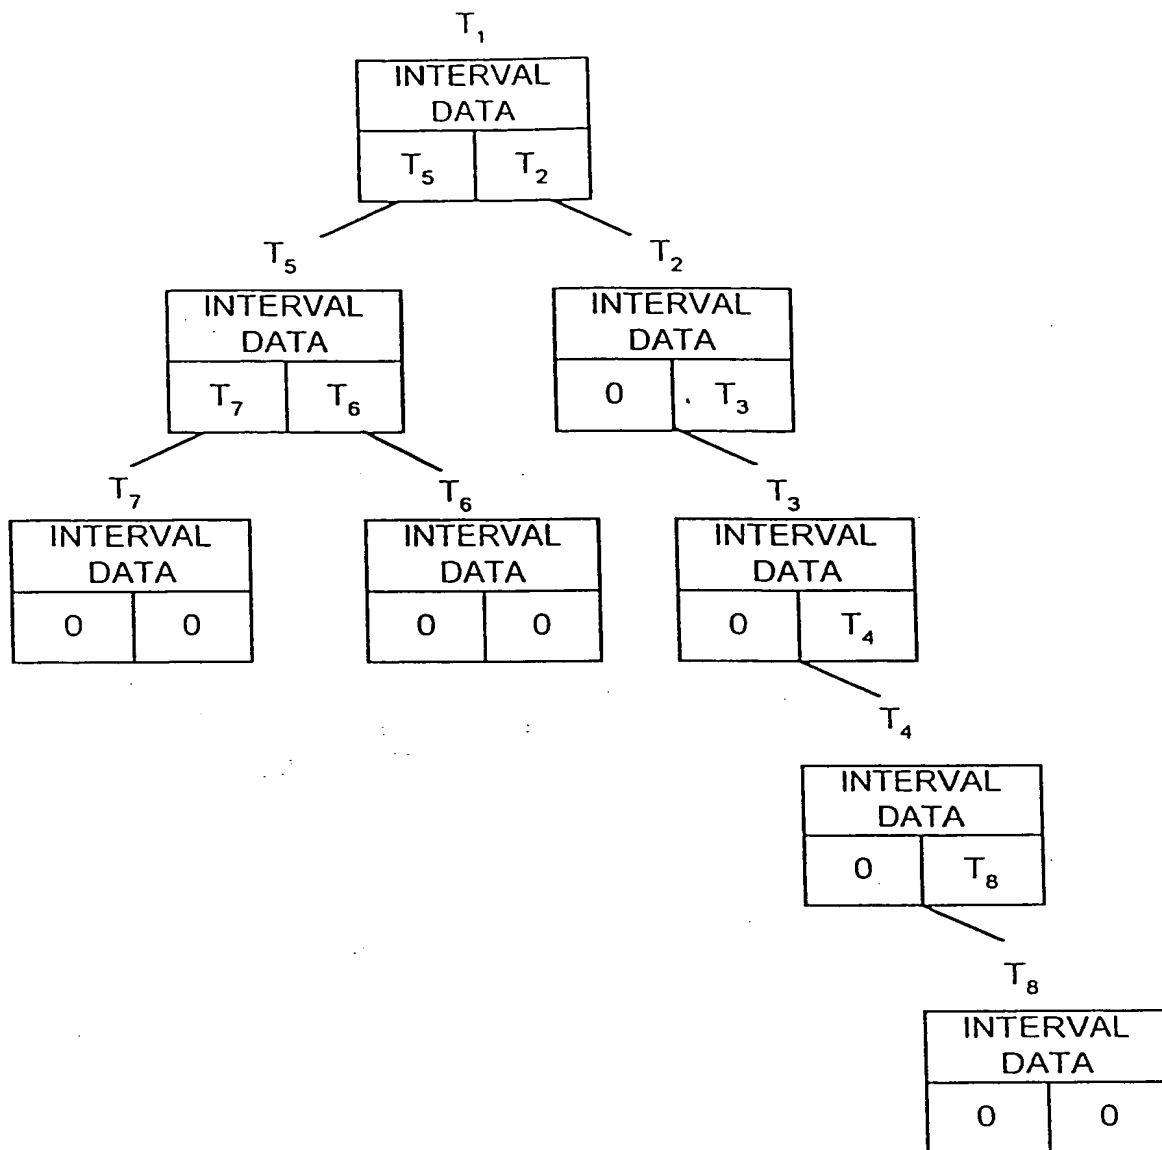



FIG. 9

This Page Blank (uspto)

FIG. 10



This Page Blank (uspto)




FIG. 11

This Page Blank (uspto)




FIG. 12

This Page Blank (uspto)




FIG. 13

This Page Blank (uspto)




FIG. 14

This Page Blank (uspto)

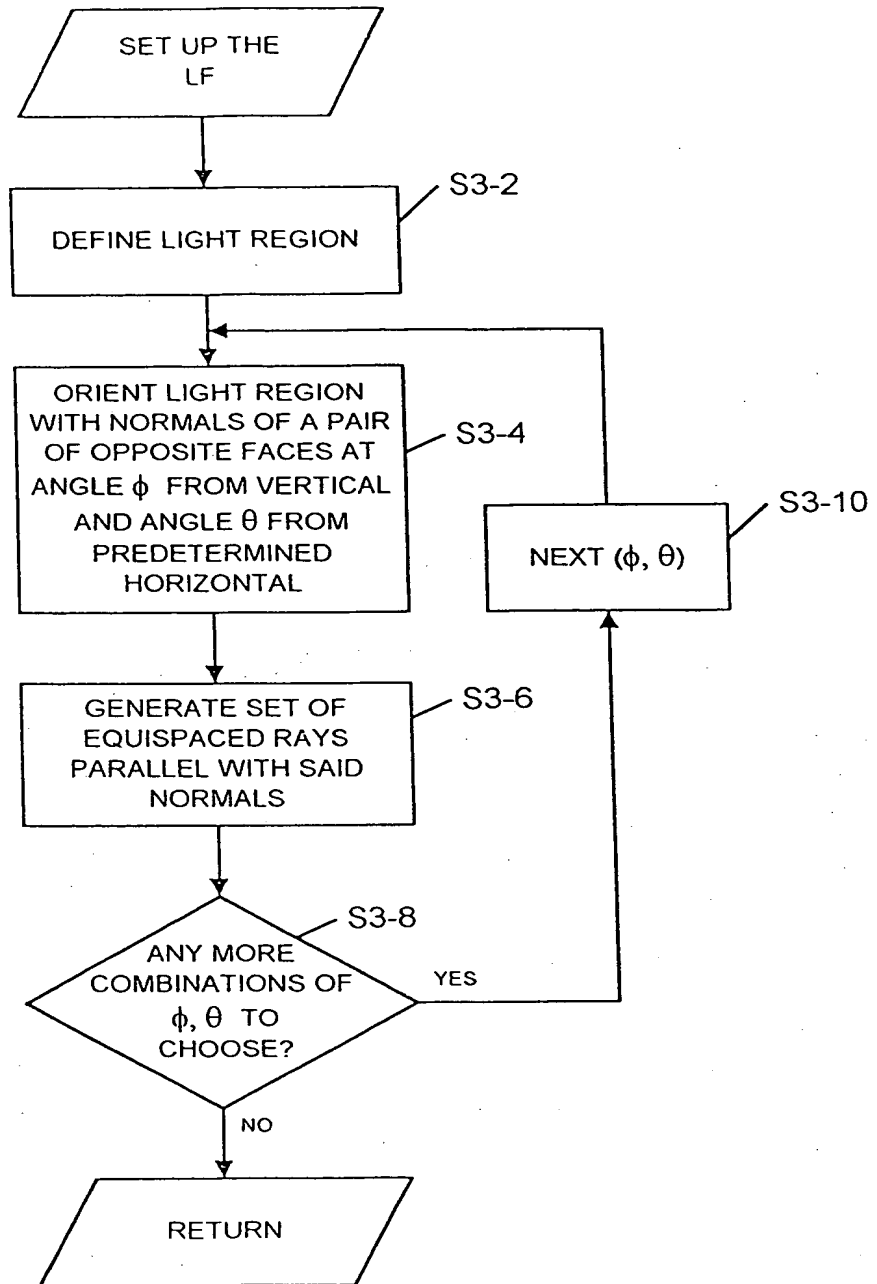


FIG. 15

This Page Blank (uspto)

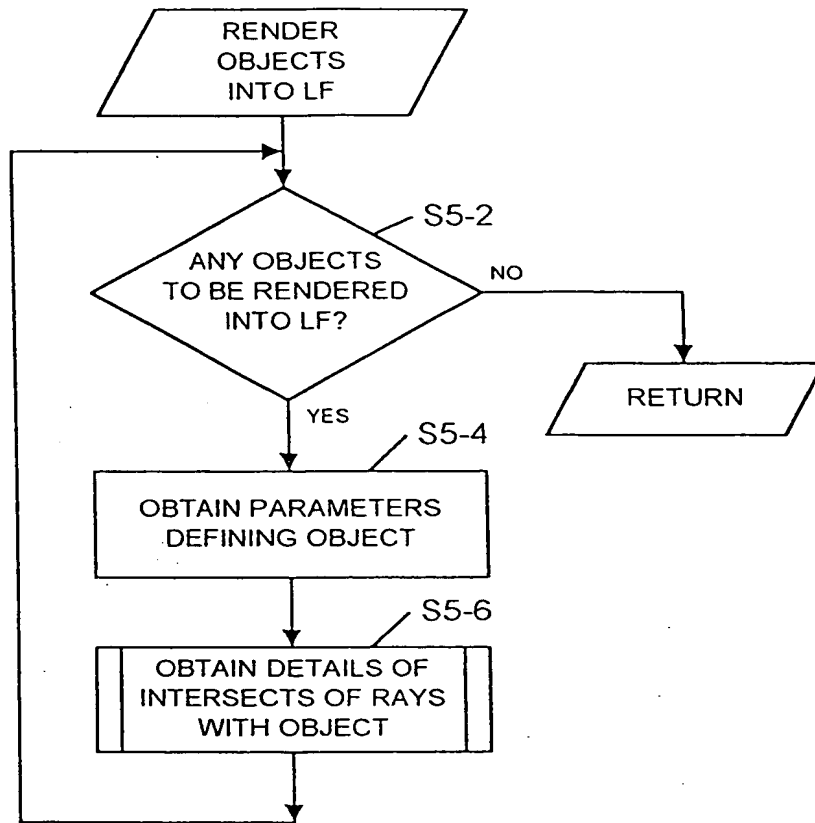


FIG. 16

This Page Blank (uspto)

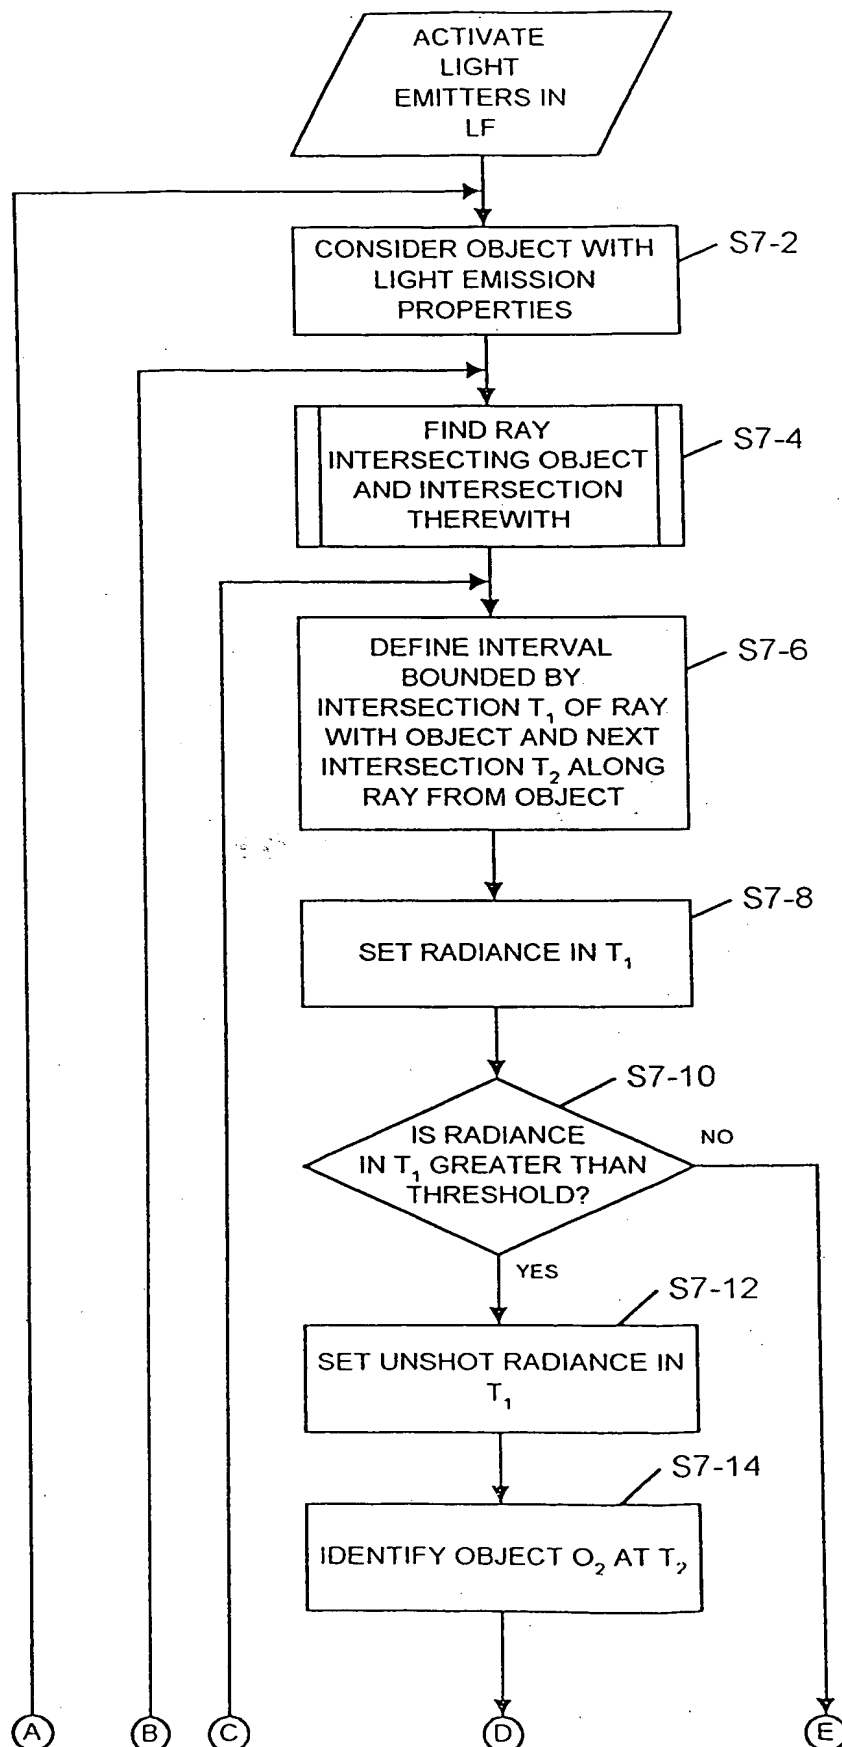


FIG. 17

This Page Blank (uspto)

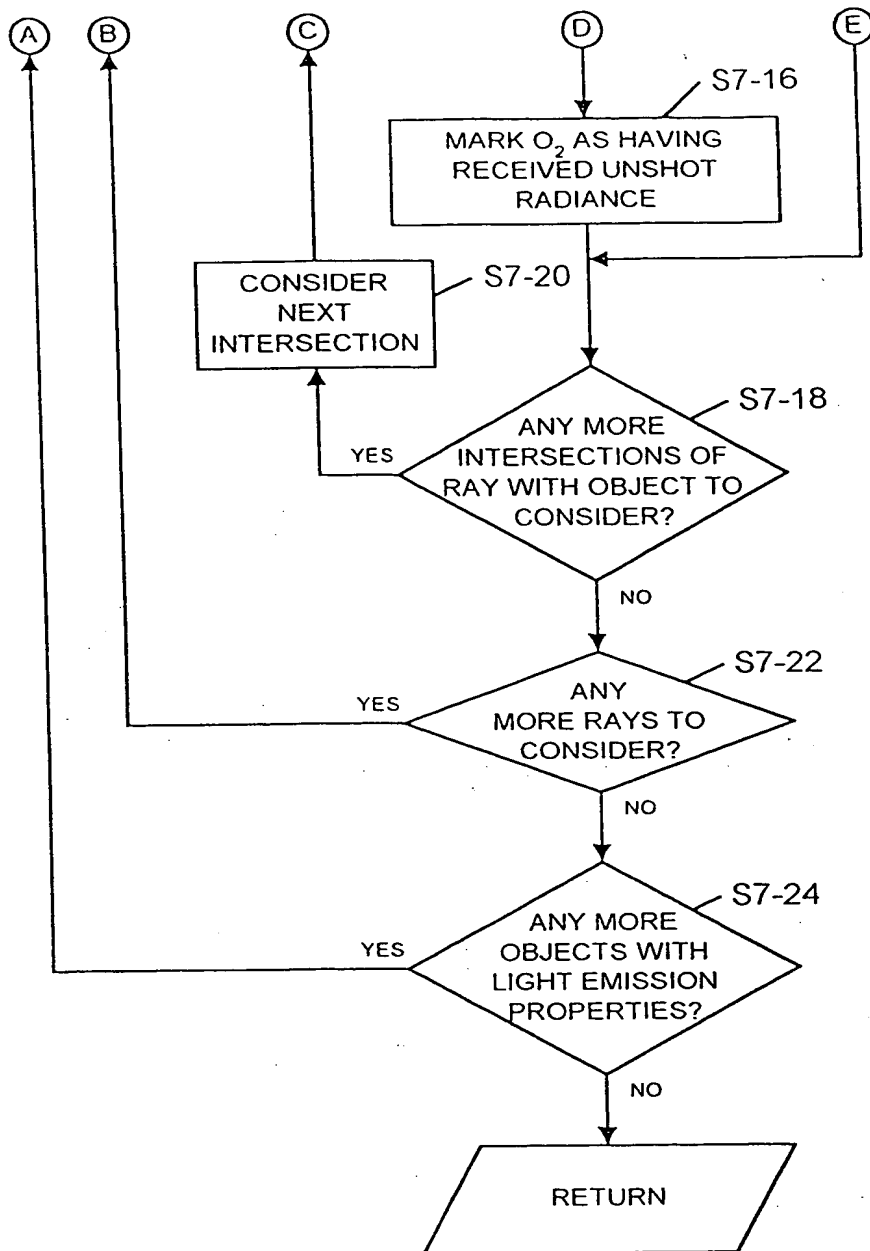


FIG. 17 (cont)

This Page Blank (uspto)

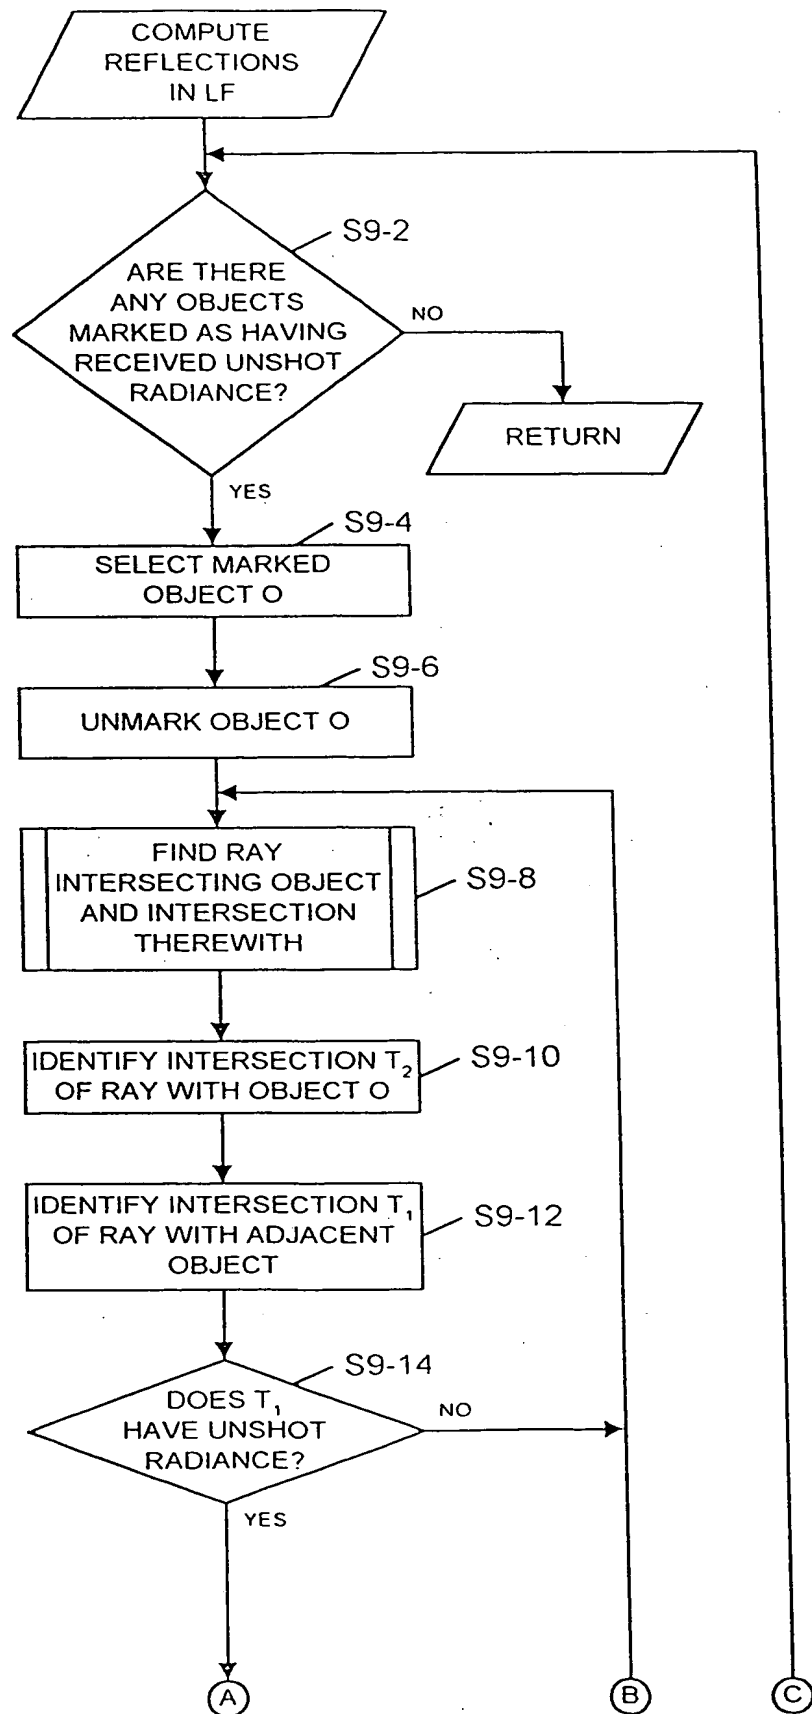


FIG. 18

This Page Blank (uspto)

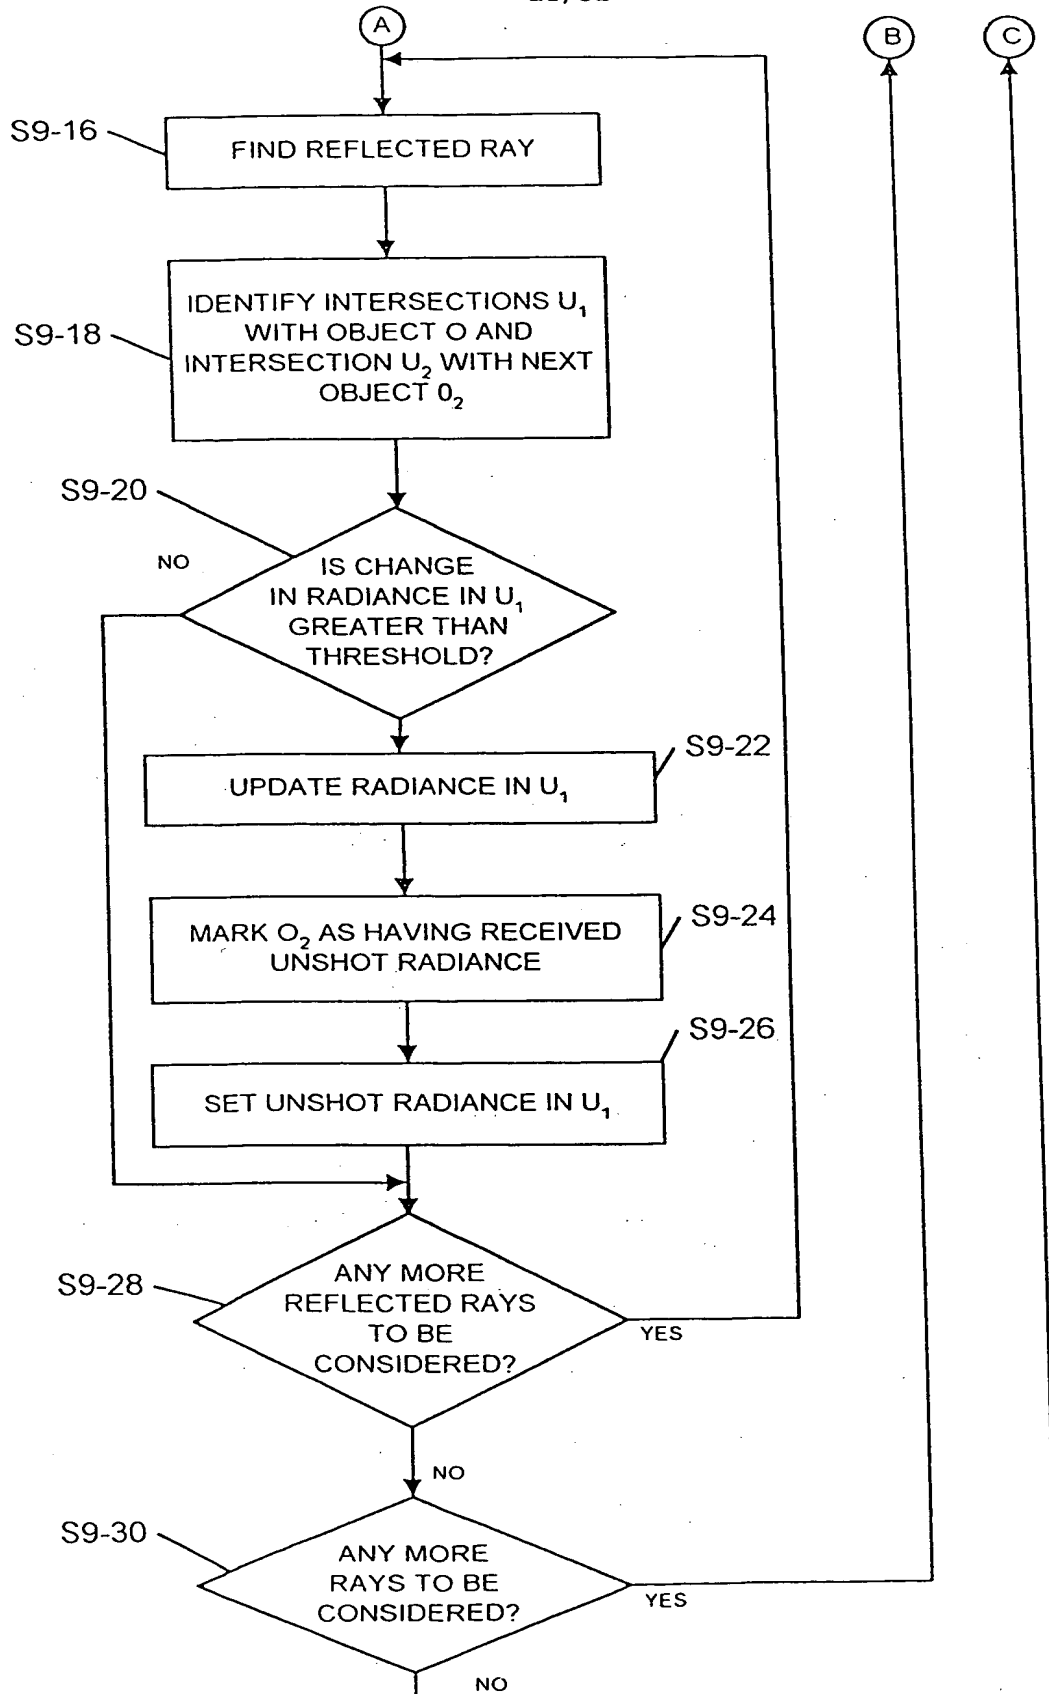


FIG. 18 (cont)

This Page Blank (uspto)

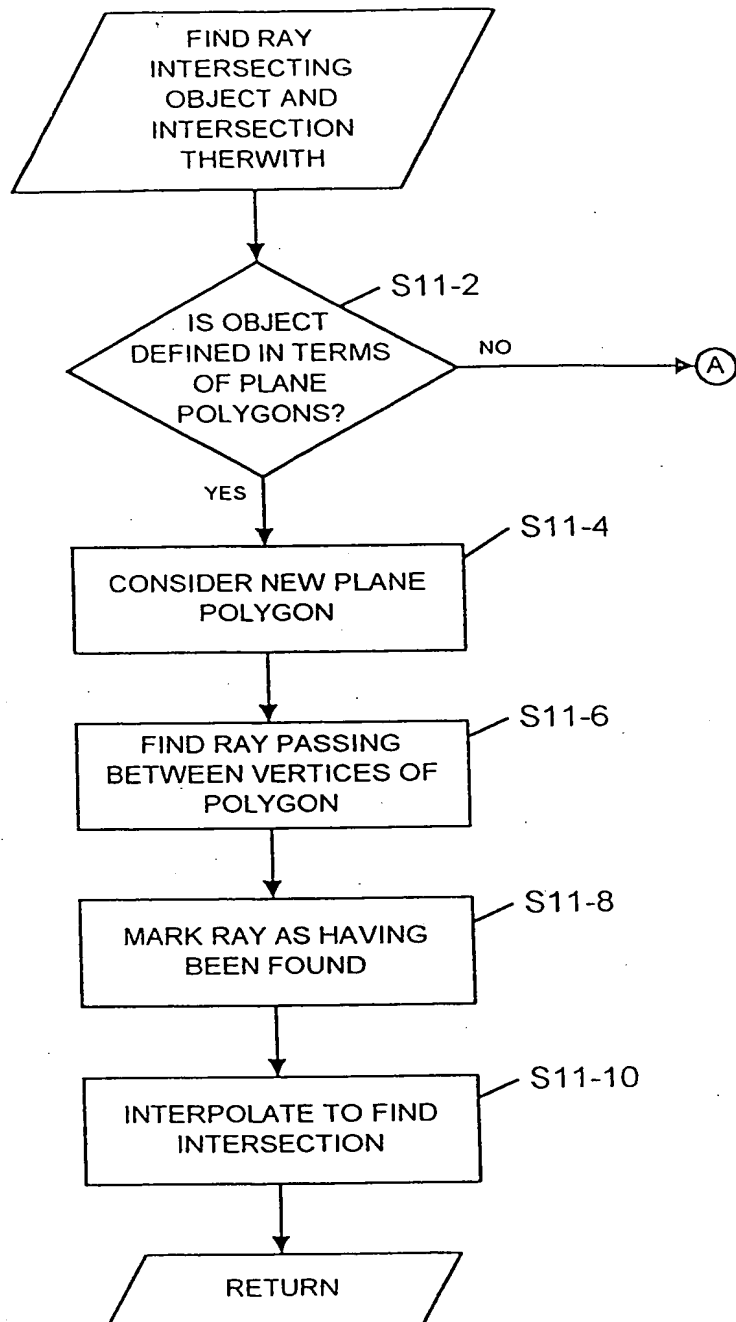


FIG. 19

This Page Blank (uspto)

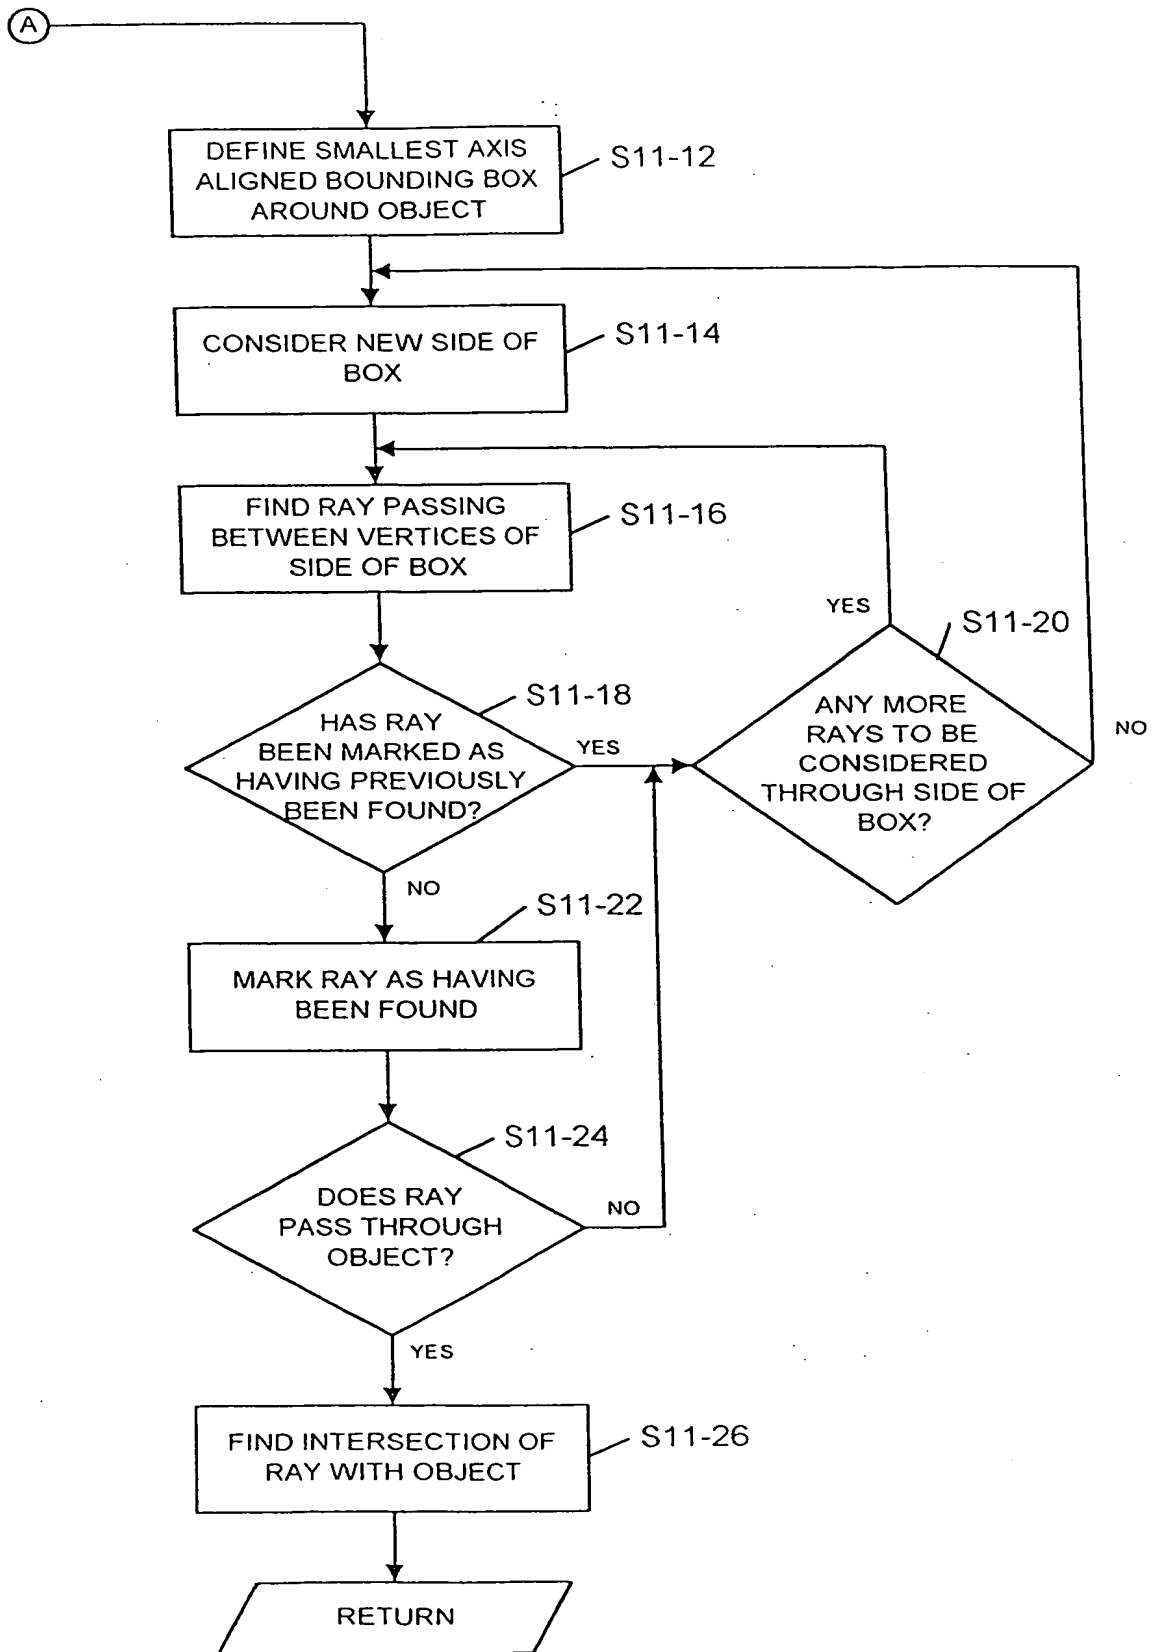


FIG. 19 (cont)

This Page Blank (uspto)

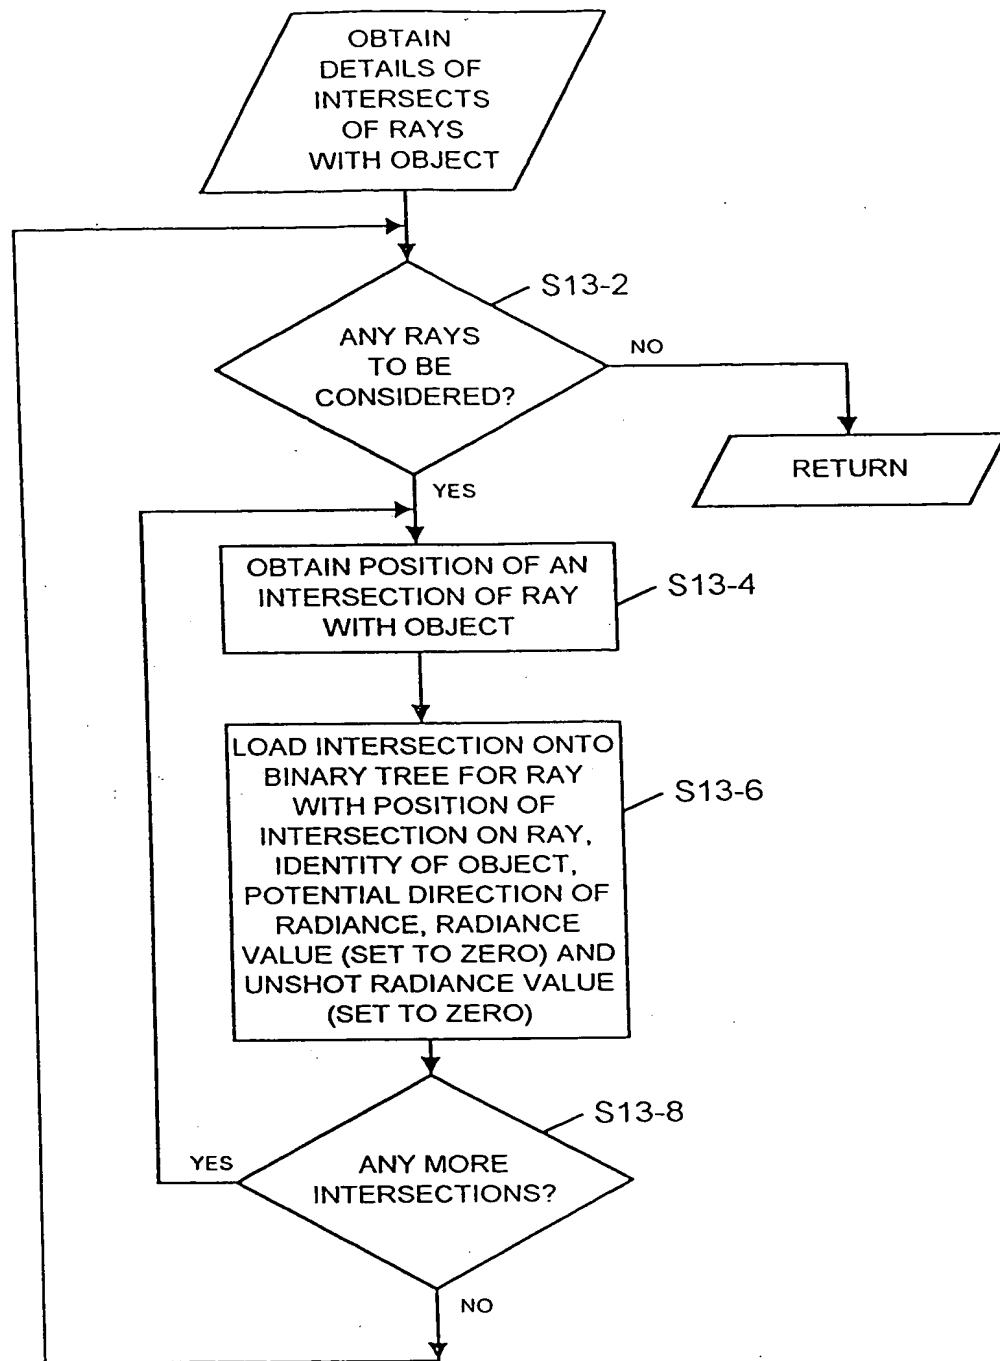


FIG. 20

This Page Blank (uspto)




FIG. 21

This Page Blank (uspto)

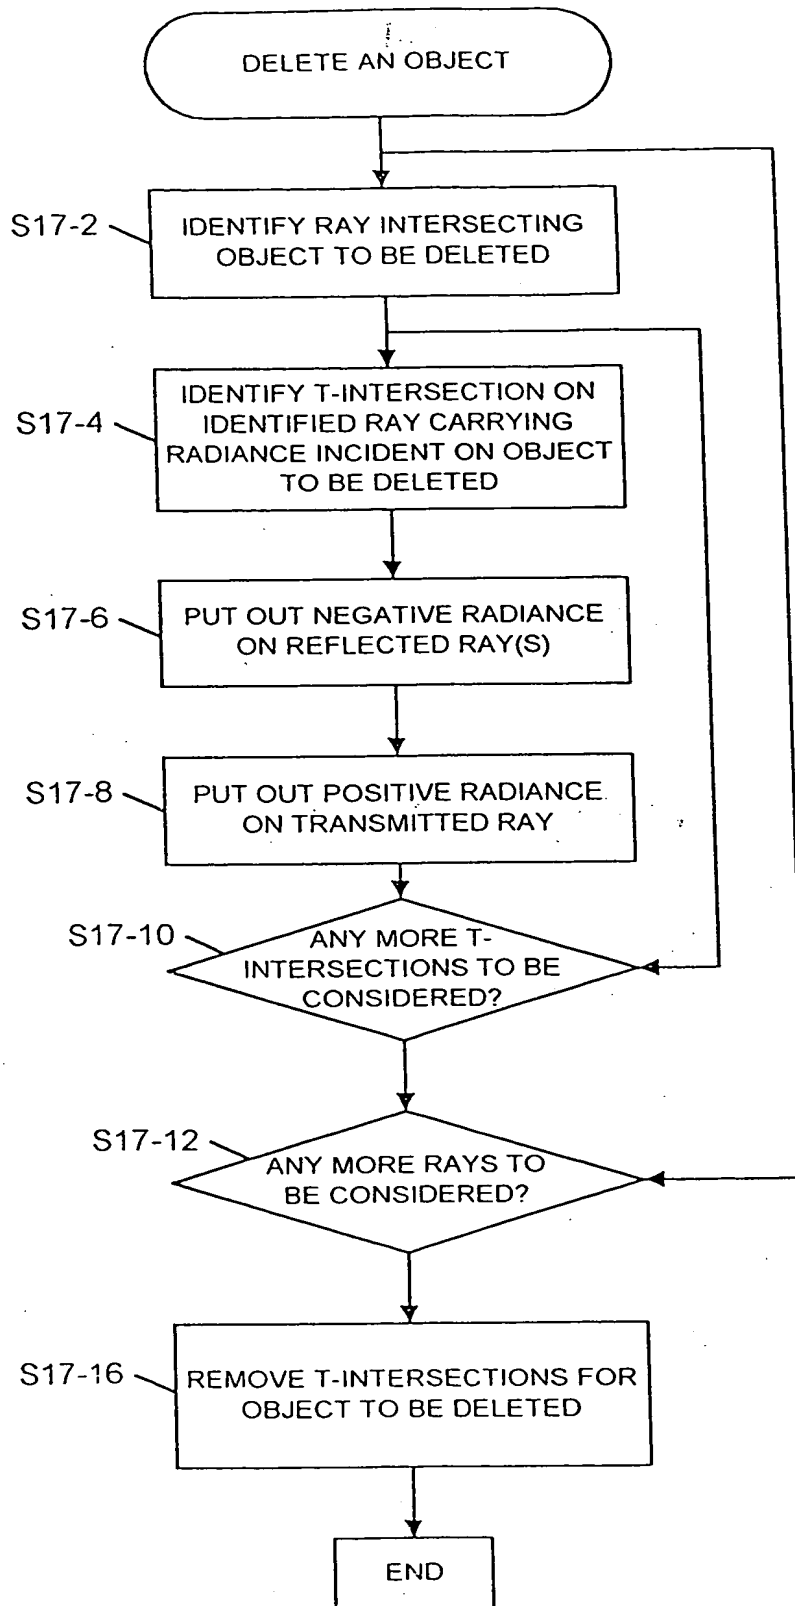


FIG. 22

This Page Blank (uspto)

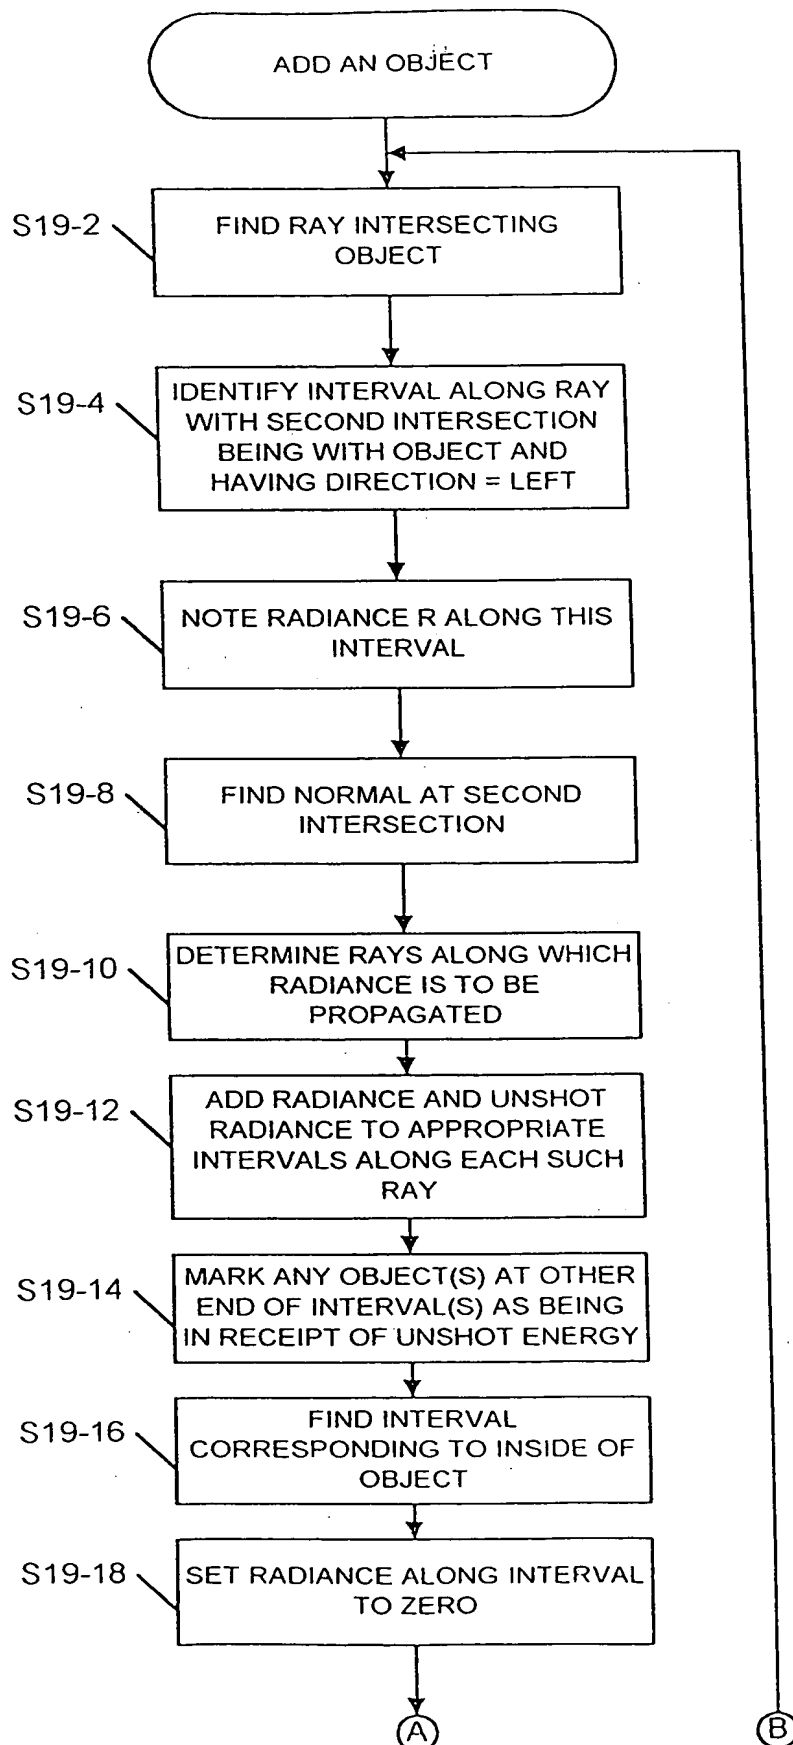


FIG. 23

This Page Blank (uspto)

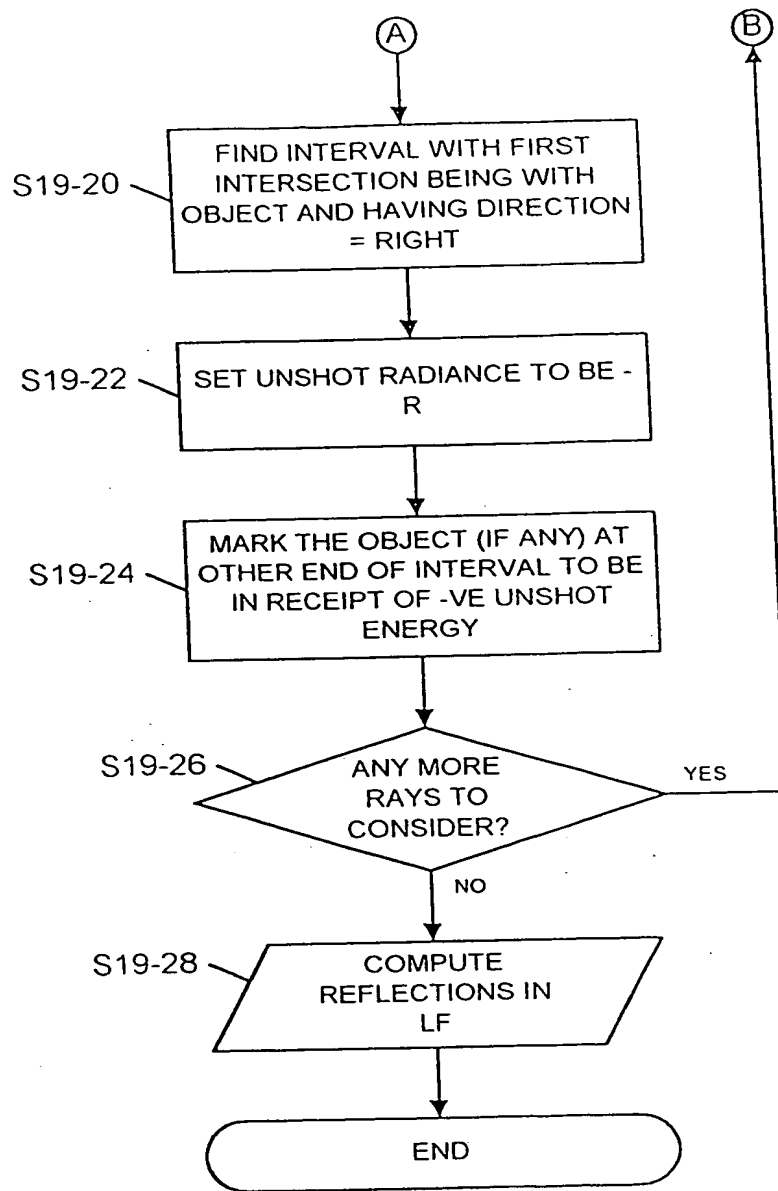


FIG. 23 (cont)

This Page Blank (uspto)

FIG. 24





FIG. 25



This Page Blank (uspto)

FIG. 26





FIG. 27



This Page Blank (uspto)




FIG. 28

This Page Blank (uspto)




FIG. 29




FIG. 30

PCT/GB 99/03229

BERESFORD + CO

23/77 29.9.99.

This Page Blank (uspto)